

The Dynamics of Learning: Excess Capacity, Interpolation, and Cognitive Structure

Flyxion

Independent Researcher

May 8, 2026

Abstract

For more than a century, cognitive science and statistical learning theory have largely operated under a shared assumption: successful learning requires compression. From Ernst Mach’s “economy of thought” to the classical bias–variance tradeoff, the dominant paradigm has treated memorization and generalization as fundamentally antagonistic processes. Under this framework, intelligence emerges by discarding the noisy particularities of experience in order to extract simplified and reusable abstractions. However, recent developments in overparameterized machine learning systems, including double descent, benign overfitting, grokking, and neural collapse, challenge the universality of this assumption. Increasing evidence suggests that systems with excess representational capacity can simultaneously memorize detailed experiences while maintaining robust generalization performance.

This essay develops a computational-theoretical framework for understanding cognition and learning under conditions of representational abundance. Central to the framework is the distinction between expressivity and representational capacity. Expressivity refers to the total representational flexibility available to a system, while capacity is defined relationally as the ratio between representational flexibility and experiential complexity. The resulting quantity determines the operational regime of the learner. Constrained regimes emerge when experiential variability exceeds available representational flexibility, forcing compression and restriction biases. Sufficient regimes arise near the interpolation threshold, where systems become maximally brittle and highly sensitive to noise. Excess regimes occur when representational flexibility substantially exceeds experiential complexity, allowing systems

to localize irregularities while preserving globally smooth representational geometry.

The framework proposes that memorization and generalization are not necessarily opposing processes. Instead, generalization may emerge through interpolation itself when excess capacity allows systems to accommodate local noise without globally destabilizing learned structure. In this view, abstraction is no longer understood primarily as compression, but as the emergence of smoothness and stability within richly overparameterized representational spaces. The essay synthesizes insights from statistical learning theory, cognitive science, developmental psychology, neuroscience, and artificial intelligence to argue that representational abundance constitutes a distinct and under-theorized regime of intelligence.

The framework additionally reinterprets developmental U-shaped learning, episodic memory, prototype formation, perceptual deterioration, and certain clinical dissociations as regime-dependent consequences of changing ratios between representational flexibility and experiential load. The essay further proposes that modern large-scale artificial systems and biological cognition may both operate through adaptive navigation across constrained, sufficient, and excess representational regimes. Finally, the work outlines the boundary conditions under which excess capacity succeeds or fails, emphasizing that compression remains adaptive in adversarial, metabolically constrained, or high-noise environments.

The central thesis is that intelligence should not be understood exclusively as a process of simplification. Under conditions of representational abundance, systems may achieve robust generalization precisely because they preserve and organize large quantities of experiential detail. Complexity, when guided by appropriate soft inductive biases, becomes a source of stability rather than error.

Contents

1	Introduction	9
2	Definitions and Scope	11
3	Historical Compression Paradigms	15
4	Representational Capacity and Experiential Complexity	18
5	The Three Regimes of Representational Capacity	23
5.1	Constrained Capacity Regimes	24
5.2	Sufficient Capacity Regimes	25
5.3	Excess Capacity Regimes	27
5.4	Representational Flexibility and High-Dimensional Separation . . .	28
5.5	The Emergence of Smoothness	29
5.6	Regime Transitions and Dynamic Cognition	30
6	Double Descent and Benign Overfitting	31
6.1	The Classical Bias–Variance Paradigm	32
6.2	The Interpolation Threshold	33
6.3	The Second Descent	34
6.4	Polynomial Interpolation and Geometric Intuition	35
6.5	Benign Overfitting	36
6.6	Minimum-Norm Interpolation	37
6.7	Volume Effects and High-Dimensional Geometry	38
6.8	The Philosophical Inversion	39
7	Smooth Interpolation and Soft Inductive Bias	40
7.1	Restriction Biases and Soft Biases	40
7.2	Interpolation Geometry	41
7.3	Minimum-Norm Dynamics	42
7.4	Flat Minima and Volume Effects	43
7.5	Localized Residual Accommodation	44
7.6	Representational Expansion and Interference Avoidance	45
7.7	The Emergence of Abstraction	46
7.8	Biological Implications	47
8	Human Empirical Signatures of Excess Capacity	47
8.1	The Classical Compression Prediction	48

8.2	Massive Recognition Fidelity	49
8.3	Superficial Detail Sensitivity	50
8.4	Prototype Effects and Emergent Smoothness	51
8.5	Perceptual Deterioration and Regime Transition	52
8.6	False Recognition and Representational Blurring	53
8.7	Noise Sensitivity and Adversarial Fragility	54
8.8	Developmental U-Shaped Learning	55
8.9	Attention, Constraint, and Gradient Preservation	56
8.10	Toward a Unified Empirical Interpretation	57
9	Developmental Dynamics and Regime Transitions	58
9.1	Development as Representational Rebalancing	59
9.2	The Geometry of U-Shaped Development	60
9.3	Constraint Emergence Through Experiential Saturation	61
9.4	Play, Constraint, and Generative Differentiation	62
9.5	The Development of Heuristics	63
9.6	False Closure and Premature Compression	64
9.7	Developmental Scaling Laws	65
9.8	Educational Implications	65
9.9	Toward Dynamic Cognitive Architecture	66
10	Clinical Dissociations and Pathological Regime Dynamics	67
10.1	Amnesia as Extreme Constraint	68
10.2	Parkinsonian Rigidity and Near-Threshold Dynamics	69
10.3	Autism and Hyper-Differentiated Representation	70
10.4	Schizophrenia and Excessive Accommodation	71
10.5	Obsessive Cognition and Recursive Constraint Locking	72
10.6	Attention Disorders and Gradient Collapse	72
10.7	Depression and Manifold Contraction	73
10.8	Clinical Phenomena as Regime Distortion	74
10.9	The One-System Hypothesis Revisited	74
11	Scaling Laws, Artificial Intelligence, and Relative Capacity	76
11.1	Artificial Neural Networks as Excess-Capacity Systems	77
11.2	Memorization Without Collapse	78
11.3	Large Language Models and Relative Constraint	79
11.4	Transformers as Dynamic Interpolation Geometry	80
11.5	Grokking and Delayed Manifold Organization	81

11.6	Adversarial Examples and Catastrophic Accommodation	82
11.7	Scaling Laws and Phase Transitions	83
11.8	The Labyrinth Problem	84
11.9	Toward Relational Artificial Intelligence	85
12	Boundary Conditions, Failure Modes, and the Cost of Excess	86
12.1	The Ecological Dependence of Learning Regimes	87
12.2	Spurious Correlations and Local Overcommitment	88
12.3	Noise Thresholds and Geometric Breakdown	89
12.4	Adversarial Environments and Defensive Compression	90
12.5	Energy Costs and Metabolic Constraints	91
12.6	The Labyrinth Problem Revisited	92
12.7	Premature Compression and Educational Failure	93
12.8	Overabundance and Semantic Diffusion	94
12.9	Constraint as Dynamic Guidance	94
12.10	Toward a Regime Ecology of Cognition	95
13	Toward a Unified Theory of Representational Abundance	96
13.1	From Compression to Geometric Organization	96
13.2	The Inversion of Memorization and Abstraction	97
13.3	The Geometry of Meaning	98
13.4	Constraint and Freedom	99
13.5	Cognition as State-Space Navigation	100
13.6	Intelligence as Interpolation Stability	100
13.7	The Philosophical Consequences	101
13.8	The End of the Economy of Thought	102
13.9	Toward a Regime Theory of Mind	102
13.10	The New Question	103
14	Conclusion	104
A	Geometric Formalization of Representational Capacity	110
A.1	Representational Systems as Maps Between Manifolds	110
A.2	Experiential Complexity	111
A.3	Representational Flexibility	112
A.4	The Operational Capacity Ratio	113
A.5	Interpolation Geometry	113
A.6	Localized Residual Accommodation Structures	115

A.7	Global Smoothness Functionals	116
A.8	Minimum-Norm Interpolation	116
A.9	Volume Effects in High Dimensions	117
A.10	Dynamic Regime Evolution	117
A.11	Generalization Through Interpolation	118
B	Statistical Learning Theory, Double Descent, and Benign Overfitting	119
B.1	Classical Bias–Variance Theory	119
B.2	The Interpolation Threshold	121
B.3	The Double Descent Transition	122
B.4	Linear Regression in the Overparameterized Regime	123
B.5	Benign Overfitting	124
B.6	Localized Accommodation Decomposition	124
B.7	Kernel Regression and Infinite-Dimensional Interpolation	125
B.8	Neural Tangent Kernels and Deep Networks	126
B.9	Flat Minima and Volume Concentration	126
B.10	Grokking as Delayed Curvature Minimization	127
B.11	Noise Scaling and Generalization Bounds	128
B.12	The Statistical Interpretation of Regime Dynamics	129
C	Dynamical Systems, Information Geometry, and Manifold Flows	130
C.1	Representational State-Spaces	130
C.2	Learning as Gradient Flow	130
C.3	Information Geometry and Fisher Metrics	131
C.4	Geodesic Interpolation	132
C.5	Curvature and Cognitive Rigidity	133
C.6	Attractor Dynamics	134
C.7	Entropy and Trajectory Volume	134
C.8	Attention as Curvature Navigation	135
C.9	Prediction Error and Manifold Correction	136
C.10	Catastrophic Overfitting as Curvature Explosion	137
C.11	Neural Collapse and Symmetry Formation	137
C.12	Representational Phase Transitions	138
C.13	The Dynamical Interpretation of Intelligence	139
D	Functional Analysis, Kernel Operators, and Spectral Geometry	140
D.1	Representational Operators	140
D.2	Kernel Representations	141

D.3	Localized Accommodation as Spectral Perturbation	141
D.4	Spectral Bias	142
D.5	Operator Smoothness and Sobolev Norms	143
D.6	Mercer Decomposition and Cognitive Modes	144
D.7	Compression as Spectral Truncation	144
D.8	The Spectral Interpretation of Overfitting	145
D.9	Attention as Operator Reweighting	145
D.10	Memory as Operator Persistence	146
D.11	The Spectral Geometry of Expertise	147
D.12	Operator Phase Transitions	147
D.13	Toward a Spectral Theory of Cognition	148
E	Category Theory, Sheaf Structures, and Cognitive Gluing	148
E.1	Representations as Morphisms	149
E.2	Representational Functors	149
E.3	Local Representational Neighborhoods	150
E.4	Gluing Conditions	151
E.5	Compression as Forced Globalization	151
E.6	Localized Accommodation and Cohomological Obstruction	152
E.7	Excess Capacity and Local Resolution	152
E.8	Representational Fiber Bundles	153
E.9	Attention as Local Trivialization	154
E.10	Functorial Generalization	154
E.11	Natural Transformations and Learning Dynamics	155
E.12	Hierarchical Sheaf Structure	155
E.13	The One-System Hypothesis Revisited Categorically	156
E.14	Toward a Sheaf-Theoretic Theory of Cognition	156
F	Thermodynamics, Entropy, and Representational Phase Space	157
F.1	Representational Phase Spaces	157
F.2	Representational Entropy	158
F.3	Experiential Entropy and Capacity	158
F.4	Free Energy and Predictive Stabilization	159
F.5	Representational Temperature	160
F.6	Phase Transitions Between Regimes	161
F.7	Localized Accommodation as Entropic Isolation	162
F.8	Catastrophic Overfitting as Entropic Propagation	162
F.9	Memory as Entropy Stabilization	163

F.10	Attention and Entropic Gradients	163
F.11	Creativity and Metastable Exploration	164
F.12	Biological Scaling and Energetic Tradeoffs	164
F.13	Toward a Thermodynamic Theory of Cognition	165
G	Variational Principles and the Euler–Lagrange Formulation of Representational Dynamics	166
G.1	Representational Action Functionals	166
G.2	The Representational Lagrangian	167
G.3	Euler–Lagrange Dynamics	167
G.4	Interpolation Constraints as Boundary Conditions	168
G.5	Localized Accommodation Fields	168
G.6	Representational Curvature Minimization	169
G.7	Noether Symmetries and Cognitive Invariants	170
G.8	Attention as Variational Focusing	170
G.9	Learning as Geodesic Flow	171
G.10	Constraint Potentials and Cognitive Rigidity	171
G.11	Entropy-Regularized Action Principles	172
G.12	Regime Dynamics as Bifurcation Theory	172
G.13	The Variational Interpretation of Intelligence	173
H	Topological Structure, Persistent Homology, and Cognitive Phase Geometry	174
H.1	Representational Topology	174
H.2	Representational Simplicial Complexes	175
H.3	Homology Groups and Cognitive Structure	176
H.4	Persistent Homology	176
H.5	Compression as Topological Collapse	177
H.6	Localized Accommodation and Topological Isolation	178
H.7	The Topology of Generalization	178
H.8	Topological Phase Transitions	179
H.9	Cognitive Rigidity as Topological Freezing	179
H.10	Semantic Diffusion as Topological Fragmentation	180
H.11	Attention as Topological Navigation	180
H.12	Memory and Topological Persistence	181
H.13	The Topological Interpretation of Abstraction	181
H.14	Toward a Topological Theory of Cognition	182

1 Introduction

The modern study of cognition emerged under the shadow of scarcity. Biological organisms possess finite metabolic resources, limited neural tissue, and incomplete access to environmental information. From these constraints arose one of the central assumptions of twentieth-century cognitive science: intelligence is fundamentally an act of compression. To learn was presumed to mean reducing the overwhelming complexity of experience into simplified, stable, and reusable abstractions. Noise, variance, and idiosyncratic detail were treated as obstacles to robust inference. Generalization depended upon their removal.

This assumption became deeply embedded across multiple intellectual traditions. In psychology, Gestalt theorists argued that perceptual organization naturally tends toward simplicity and regularity. In philosophy of science, Ernst Mach proposed the “economy of thought,” according to which cognition seeks compressed conceptual structures that minimize unnecessary complexity. In statistics and machine learning, the classical bias–variance tradeoff formalized the intuition that increasing model complexity eventually leads to catastrophic overfitting. Across these traditions, the underlying principle remained remarkably stable: successful intelligence requires abstraction through reduction.

Yet contemporary developments in machine learning have destabilized this foundational assumption. Overparameterized systems frequently violate the traditional relationship between complexity and generalization. Deep neural networks containing vastly more parameters than training examples are capable of perfectly interpolating noisy datasets while still generalizing effectively to unseen data. Phenomena such as double descent, benign overfitting, grokking, and neural collapse suggest that the relationship between memorization and abstraction may not be fundamentally antagonistic. Under certain conditions, systems appear capable of simultaneously preserving fine-grained experiential detail while maintaining globally stable predictive structure.

These observations create a profound theoretical tension. If high-capacity systems can memorize and generalize simultaneously, then the traditional compression-centered account of cognition may only describe a restricted subset of possible learning regimes. The historical assumption that intelligence requires simplification may reflect the study of systems operating under conditions of representational scarcity rather than a universal law of learning itself.

This essay develops an alternative framework centered on the concept of repre-

sentational abundance. The framework proposes that cognitive behavior depends not on absolute complexity alone, but on the relationship between representational flexibility and experiential complexity. Under this view, learning systems occupy different operational regimes depending on the ratio between available representational resources and the diversity of situations they must encode. Compression becomes only one possible adaptive strategy among several.

The framework distinguishes between three primary regimes. In constrained regimes, experiential complexity exceeds representational flexibility, forcing the learner to discard detail and rely on restriction biases to achieve generalization. In sufficient regimes, representational flexibility approximately matches experiential complexity, producing brittle interpolation and maximal sensitivity to noise. In excess regimes, representational flexibility substantially exceeds experiential complexity, allowing the system to preserve detailed experiential structure while maintaining globally smooth representational organization.

The central claim of this work is that generalization may emerge through interpolation rather than despite it. In excess-capacity systems, memorization and abstraction are not opposing forces. Instead, local irregularities can be accommodated within highly expressive representational spaces without globally destabilizing learned structure. Complexity becomes not a liability, but a stabilizing substrate capable of supporting both specificity and generality simultaneously.

Importantly, this framework does not reject compression entirely. Constrained representations remain adaptive under many environmental conditions, particularly when systems face severe metabolic limits, adversarial perturbations, sparse data, or highly noisy distributions. The present argument is therefore not that simplicity is obsolete, but that simplicity represents only one region within a broader geometry of learning. Compression-centered theories become regime-specific descriptions rather than universal principles.

The implications of this shift extend across multiple domains. In cognitive science, the framework offers reinterpretations of episodic memory, prototype formation, perceptual deterioration, developmental U-shaped learning, and clinical dissociations traditionally modeled using dual-process architectures. In artificial intelligence, the framework provides a conceptual bridge between overparameterized learning systems and biological cognition. In philosophy and epistemology, the framework challenges the longstanding assumption that abstraction necessarily requires the elimination of detail.

The objective of this essay is not merely to introduce a new metaphor for

cognition, but to construct a regime theory of learning itself. To achieve this, the following sections progressively develop the formal ontology, computational mechanisms, empirical signatures, developmental consequences, and theoretical implications of representational abundance. We begin by establishing the foundational definitions and scope conditions necessary for the framework.

2 Definitions and Scope

Any theoretical framework attempting to unify cognition, statistical learning, and artificial intelligence must establish a precise ontology before advancing explanatory claims. Historically, many interdisciplinary theories have failed because key terms such as “representation,” “memory,” “generalization,” and “complexity” shifted meaning across levels of analysis without formal stabilization. The present framework therefore begins by defining its primary concepts explicitly and delimiting the scope of its claims.

The framework concerns systems that learn predictive relationships between situations and outcomes. A cognitive system, in the broadest sense employed here, is any biological or artificial system capable of modifying its internal structure on the basis of prior experience in order to anticipate future conditions. This definition intentionally includes human cognition, animal learning, statistical models, and modern neural architectures under a unified computational abstraction.

A situation refers to an external context, stimulus configuration, or informational state encountered by the learner. Formally, situations may be represented as elements

$$x \in \mathcal{X},$$

where \mathcal{X} denotes the space of possible inputs or environmental conditions.

An experience is defined as a pairing between a situation and its associated outcome. Experiences therefore take the form

$$(x, y),$$

where

$$x \in \mathcal{X}, \quad y \in \mathcal{Y},$$

and \mathcal{Y} denotes the space of possible outcomes, labels, actions, or consequences.

Learning consists of constructing internal transformations capable of mapping

situations to anticipated outcomes. These transformations constitute representations. A representation is therefore not simply stored information, but an organized internal geometry that enables predictive coordination across situations.

The framework distinguishes sharply between expressivity and capacity. This distinction is foundational.

Definition 1 (Expressivity). *Expressivity is the total representational flexibility available to a system. It corresponds to the volume of internal transformations, partitions, or functions that the system can potentially realize.*

In artificial neural systems, expressivity is often related to parameter count, architectural depth, connectivity structure, or function-space complexity. In biological systems, expressivity may emerge from neural population diversity, cortical expansion, synaptic structure, recurrent connectivity, or dynamical flexibility. Importantly, expressivity is treated here as an abstract geometric property rather than a single measurable scalar quantity.

Expressivity alone, however, is insufficient for characterizing learning behavior. A system containing one billion parameters may still operate under representational scarcity if its experiential demands are sufficiently large. Conversely, a comparatively small biological subsystem may operate under representational abundance if its experiential domain remains narrow and highly structured.

The critical quantity is therefore representational capacity.

Definition 2 (Representational Capacity). *Representational capacity is the relationship between a system's effective representational flexibility and the complexity of the experiences it must encode.*

The framework introduces the relational quantity

$$\rho = \frac{C}{E},$$

where C denotes effective representational flexibility and E denotes experiential complexity or experiential load.

The quantity ρ determines the operational regime of the learner.

Importantly, neither C nor E should initially be interpreted as strictly fixed numerical quantities. At the present stage of the framework, they function as abstract geometric and informational descriptors. Later refinements may operationalize them through quantities such as manifold dimensionality, effective rank, Fisher information, mutual information, parameter norms, trajectory entropy, or repre-

sentational sparsity. However, premature commitment to any single formalization risks obscuring the broader conceptual structure the framework seeks to establish.

Experiential complexity refers not merely to the number of experiences, but to the structural diversity, variability, and informational irregularity of those experiences. Two datasets containing the same number of observations may induce radically different experiential complexity depending on redundancy, smoothness, noise structure, and latent dimensionality.

The distinction between local and global structure is equally central to the framework. Local structure refers to fine-grained accommodation of specific experiences, exceptions, or irregularities. Global structure refers to large-scale smoothness, continuity, or stable predictive organization across the representational manifold.

Classical compression-centered theories generally assume that preserving local irregularities necessarily disrupts global smoothness. The present framework challenges this assumption by proposing that sufficiently expressive systems can localize irregularities without globally destabilizing representation.

To formalize this intuition, we introduce the concept of localized residual accommodation.

Definition 3 (Localized Residual Accommodation). *A localized residual accommodation structure is a confined representational modification that enables a system to interpolate specific irregular experiences while preserving smoothness across the broader representational manifold.*

The pedagogical metaphor of “spikes” used throughout this framework refers to such localized accommodation structures. These need not correspond literally to one-dimensional polynomial spikes. Depending on implementation, they may emerge as high-frequency residual components, local curvature corrections, sparse interpolation pockets, attractor basins, kernel activations, or partitioned attention structures. The framework intentionally remains implementation-neutral at this stage.

The distinction between hard and soft inductive biases also requires careful definition.

Definition 4 (Restriction Bias). *A restriction bias is a hard architectural or representational limitation that forces simplification because the system lacks the flexibility required to interpolate all experiences simultaneously.*

Definition 5 (Soft Inductive Bias). *A soft inductive bias is an emergent optimization preference that guides the system toward specific solutions among many available interpolating configurations.*

Restriction biases dominate constrained systems. Soft inductive biases dominate excess-capacity systems.

This distinction is crucial because classical theories often conflate simplicity produced by incapacity with smoothness produced by optimization geometry. The present framework argues that these are fundamentally different mechanisms of generalization.

The scope of the framework must also be carefully bounded. The present work does not claim that all cognition operates exclusively in excess-capacity regimes. Nor does it claim that biological systems literally implement the same optimization procedures used in modern deep learning systems. Instead, the framework proposes that several apparently distinct phenomena across cognition and machine learning may reflect common geometric principles governing learning under different ratios of representational flexibility to experiential complexity.

Similarly, the framework does not deny the existence of modular specialization, functional dissociation, or heterogeneous neural dynamics. Rather, it proposes that many dual-process phenomena may emerge from different operational regimes within a unified representational substrate rather than requiring entirely separate learning principles.

Finally, the framework distinguishes carefully between empirical observations, computational reinterpretations, mechanistic hypotheses, speculative extensions, and philosophical implications. The empirical phenomena discussed throughout the essay are well-established findings from cognitive science and machine learning. The regime-theoretic interpretation proposed here constitutes a computational synthesis intended to organize these findings within a unified representational geometry.

Having established the foundational ontology and scope conditions of the framework, we may now examine the historical emergence of compression-centered theories and the assumptions that made representational scarcity appear inevitable.

3 Historical Compression Paradigms

The modern scientific understanding of cognition emerged within a conceptual landscape dominated by scarcity. Biological organisms appeared metabolically constrained, environmentally overwhelmed, and computationally limited. Under these conditions, the prevailing assumption became almost unavoidable: intelligence must function primarily through simplification. The richness of the external world could not possibly be preserved in full detail, and therefore cognition was understood as a process of selective reduction. Learning meant compressing complexity into manageable abstractions.

Although this assumption appeared across multiple intellectual traditions, its underlying logic remained remarkably consistent. Systems capable of discarding irrelevant detail were presumed more adaptive than systems attempting to preserve experiential specificity. Simplicity was therefore elevated from an engineering convenience into a normative principle of intelligence itself.

One of the earliest and most influential formulations emerged through the work of $\square\square$. Mach proposed that scientific and cognitive thought operate according to an “economy of thought,” in which the mind reduces diverse phenomena into compressed conceptual regularities. Under this view, cognition achieves efficiency by minimizing unnecessary representational complexity. The mind experiences conceptual relief when multiple phenomena become reducible to a smaller number of organizing principles.

This principle of economy became deeply influential because it aligned naturally with broader assumptions about biological limitation. If cognition is metabolically expensive, then reducing representational burden appears inherently adaptive. Simplicity therefore became associated not merely with elegance, but with survival itself.

Gestalt psychology extended this orientation into perceptual organization. The Gestalt principle of *Pragnanz* argued that perceptual systems naturally organize stimuli into the “best,” simplest, and most stable forms possible. Complex or ambiguous sensory information was presumed to collapse toward regularized structures emphasizing symmetry, continuity, and simplicity.

Importantly, the Gestalt notion of “good form” implicitly tied successful cognition to representational reduction. Simpler organizations were considered psychologically privileged because they minimized computational complexity while preserving functional coherence.

This same intuition later reappeared within formal statistical learning theory. Classical statistics treated generalization as a balancing act between underfitting and overfitting. Systems with insufficient complexity failed to capture relevant structure, producing high bias. Systems with excessive complexity became hypersensitive to noise, producing high variance. The resulting bias–variance tradeoff implied that successful learning requires locating an intermediate region of controlled simplicity.

Under this framework, memorization became synonymous with failure. A learner capable of perfectly fitting noisy training data was presumed incapable of generalizing beyond it. Interpolation and robustness appeared fundamentally incompatible.

The same compression-centered logic permeated cognitive science throughout the twentieth century. Prototype theories of categorization proposed that cognition abstracts generalized central tendencies from collections of specific experiences. Symbolic models treated intelligence as the extraction of compressed rules governing broad classes of situations. Even memory systems were frequently divided according to the distinction between abstraction and specificity. Episodic memory preserved local detail, while semantic systems extracted generalized structure.

Several influential theories formalized this division directly. Complementary Learning Systems theory argued that rapid memorization of individual episodes and slow extraction of statistical regularities required distinct computational architectures because the two objectives interfere with one another. Preserving arbitrary traces was presumed computationally incompatible with extracting stable invariances.

The compression paradigm therefore became deeply institutionalized across multiple levels of explanation. Biological limitation justified abstraction. Statistical learning theory formalized simplicity. Cognitive science partitioned memory systems according to the presumed incompatibility between memorization and generalization.

Yet despite its historical dominance, the compression paradigm always contained unresolved tensions.

Human cognition repeatedly exhibited surprisingly high-fidelity memory. Experimental work demonstrated that individuals could retain detailed recognition of thousands of visual objects with remarkable accuracy. Exemplar effects revealed

that people often rely upon highly specific experiential traces even when abstract rules are available. Developmental learning frequently exhibited non-monotonic trajectories inconsistent with simple abstraction models. Expertise sometimes increased sensitivity to subtle local distinctions rather than reducing representation toward compressed summaries.

These anomalies were frequently treated as secondary phenomena operating around a fundamentally compression-centered architecture. However, the emergence of modern overparameterized machine learning systems dramatically altered the situation. Neural networks capable of perfectly interpolating massive datasets nevertheless generalized successfully. Systems containing vastly more parameters than training examples violated the classical expectation that interpolation necessarily destroys robustness.

The discovery of double descent represented a particularly significant rupture. Under classical assumptions, increasing complexity beyond the interpolation threshold should monotonically worsen generalization. Instead, highly overparameterized systems often exhibited a second descent in error after passing through the interpolation regime. Generalization recovered precisely where classical theory predicted catastrophic failure.

This observation destabilized one of the central assumptions underlying compression-centered cognition: that memorization and abstraction necessarily compete for representational resources.

At the same time, phenomena such as benign overfitting, grokking, neural collapse, and implicit regularization suggested that highly expressive systems possess emergent organizational dynamics not captured by traditional bias–variance intuitions. Rather than globally destabilizing representation, excess expressivity sometimes enabled systems to localize irregularity while preserving large-scale smoothness.

These developments do not invalidate the historical compression paradigm entirely. Under many environmental conditions, constrained representations remain adaptive. Simplicity often provides robustness under noise, metabolic limitation, sparse data, adversarial perturbation, or rapidly shifting environments. Compression-centered cognition therefore remains an important and legitimate operating regime.

However, the historical mistake may have been treating constrained learning as the universal form of intelligence rather than one region within a broader

representational geometry.

The framework developed in this essay proposes that the classical compression paradigm primarily describes systems operating under conditions where representational flexibility remains limited relative to experiential complexity. Under such conditions, simplification becomes unavoidable because the learner lacks sufficient representational freedom to preserve experiential detail.

But when representational flexibility substantially exceeds experiential complexity, a fundamentally different regime emerges. In this regime, systems need not choose between memorization and generalization. They may instead preserve detailed experiential structure while maintaining globally smooth predictive organization.

The transition from compression-centered cognition to representational abundance therefore represents more than a technical adjustment within machine learning theory. It represents a shift in the ontology of intelligence itself.

To formalize this shift rigorously, we must now develop a mathematical and geometric account of representational capacity and the operational regimes that emerge from its relationship to experiential complexity.

4 Representational Capacity and Experiential Complexity

The central claim of the present framework is that learning behavior cannot be understood through model complexity alone. Neither parameter count, neural volume, architectural depth, nor computational power independently determines whether a system memorizes, generalizes, compresses, or interpolates. These properties emerge relationally through the interaction between representational flexibility and experiential complexity.

Historically, discussions of intelligence frequently treated capacity as an absolute quantity. A system was assumed to possess either “high” or “low” capacity independently of the structure of the environment it encountered. However, this assumption collapses under closer examination. A system containing billions of parameters may still operate under representational scarcity if exposed to sufficiently diverse and irregular experiential distributions. Conversely, a comparatively small biological subsystem may operate under representational abundance if its experiential manifold remains narrow and structured.

The relevant quantity is therefore not absolute size, but relative representational freedom.

We define the operational capacity ratio as

$$\rho = \frac{C}{E},$$

where C denotes effective representational flexibility and E denotes experiential complexity.

The quantity ρ governs the operational learning regime of the system.

The numerator C should not be interpreted merely as parameter count. Effective representational flexibility depends on the geometry of the representational manifold accessible to the learner. In artificial systems, this may depend upon parameterization, connectivity, optimization dynamics, activation structure, depth, attention mechanisms, or feature reuse. In biological systems, it may depend upon cortical expansion, recurrent architecture, neural diversity, temporal dynamics, synaptic plasticity, or population coding structure.

Similarly, experiential complexity E is not reducible to the number of observations alone. Two environments containing identical numbers of experiences may impose radically different representational demands depending upon variability, latent dimensionality, redundancy structure, noise distribution, or topological irregularity.

The framework therefore treats both C and E as effective geometric quantities rather than simple scalar measurements.

To formalize this more precisely, consider a learning system attempting to approximate a mapping

$$f : \mathcal{X} \rightarrow \mathcal{Y},$$

where \mathcal{X} denotes the space of situations and \mathcal{Y} denotes the space of anticipated outcomes.

The learner constructs an internal representation

$$\phi : \mathcal{X} \rightarrow \mathcal{Z},$$

where \mathcal{Z} denotes an internal representational manifold.

Classical compression-centered theories implicitly assume that successful learning requires reducing the dimensionality or complexity of \mathcal{Z} relative to \mathcal{X} . Under

this assumption, robust generalization emerges by collapsing many distinct situations into simplified equivalence classes.

The present framework proposes a more general possibility. In excess-capacity systems, the dimensionality and geometric flexibility of \mathcal{Z} may substantially exceed the minimal structure required to interpolate the observed distribution. Rather than compressing experiential detail aggressively, the system may preserve large quantities of local variation while organizing those variations within globally smooth manifolds.

This distinction is subtle but fundamental.

Compression-centered systems generalize because they are forced to ignore detail.

Excess-capacity systems generalize because they possess enough representational flexibility to localize detail without globally destabilizing representation.

The difference between these mechanisms becomes particularly clear when examining interpolation.

Suppose a dataset

$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$$

contains noisy observations sampled from an underlying latent process. Classical learning theory predicts that systems capable of perfectly fitting all observations will necessarily overfit noise and generalize poorly.

However, modern overparameterized systems frequently violate this prediction.

To understand why, consider the geometry of interpolation itself.

A constrained system lacks sufficient representational flexibility to satisfy all interpolation constraints simultaneously. Formally, if the effective representational dimensionality is smaller than the complexity of the experiential distribution, then exact interpolation becomes impossible without aggressive simplification. The learner must therefore discard variance and emphasize low-dimensional summaries.

Near the interpolation threshold, however, the system acquires just enough flexibility to satisfy all interpolation constraints exactly. This regime corresponds to sufficient capacity.

At sufficient capacity, the learner possesses no surplus representational free-

dom beyond what is required for interpolation itself. Consequently, interpolation becomes globally unstable. Small fluctuations in individual observations force large-scale distortions throughout the representational manifold.

This instability can be illustrated geometrically. Let the learner seek a function f_θ parameterized by θ satisfying

$$f_\theta(x_i) = y_i$$

for all training examples.

Near the interpolation threshold, the solution space becomes extremely constrained. To satisfy every interpolation condition simultaneously, the learner must construct highly irregular global deformations. The resulting solution exhibits extreme sensitivity to local perturbation.

Classical overfitting emerges naturally from this regime.

The situation changes fundamentally once representational flexibility substantially exceeds experiential complexity.

When

$$\rho \gg 1,$$

the system possesses many more representational degrees of freedom than are strictly necessary for interpolation. The solution space therefore expands dramatically.

Instead of requiring global distortions to satisfy local irregularities, the learner may localize accommodation within confined representational subregions while preserving smoothness elsewhere.

This produces what the framework terms localized residual accommodation structures.

To illustrate the intuition, consider a decomposition

$$f_\theta(x) = g(x) + r(x),$$

where $g(x)$ captures globally smooth structure and $r(x)$ denotes localized residual accommodations required to interpolate noisy observations.

In constrained or sufficient regimes, the residual term cannot remain localized because representational freedom is insufficient to isolate local irregularities. Noise therefore propagates globally through the learned representation.

In excess regimes, however, the residual accommodations may remain geometrically confined:

$$|r(x)| \approx 0$$

for most regions of the manifold, while becoming nontrivial only near highly specific experiential neighborhoods.

The learner therefore preserves detailed interpolation without globally sacrificing smoothness.

This mechanism provides the geometric intuition underlying benign overfitting. The learner memorizes local irregularities while maintaining globally stable predictive organization.

Importantly, this framework does not claim that systems literally generate new information beyond the observed data. Rather, excess-capacity systems possess representational geometries containing more internal degrees of freedom than are minimally required for interpolation. The resulting manifolds support richer partitioning, smoother transitions, and more stable accommodation of local variance.

The distinction between local and global geometry becomes essential.

Classical learning theory often implicitly assumes that local accommodation necessarily disrupts global organization. Excess-capacity systems violate this assumption because their representational manifolds possess sufficient flexibility to decouple local irregularity from global instability.

This decoupling also clarifies the distinction between restriction biases and soft inductive biases.

In constrained systems, generalization emerges because representational freedom is limited. Simplicity is enforced structurally. The learner literally cannot represent arbitrary detail.

In excess systems, by contrast, the learner possesses infinitely many interpolating solutions. Generalization therefore depends upon optimization geometry rather than structural incapacity.

The learner selects among interpolating solutions according to emergent preferences induced by optimization dynamics.

This distinction can be formalized through minimum-norm interpolation.

Suppose the learner seeks parameters satisfying the interpolation constraints

$$X\theta = y,$$

where X denotes the design matrix and y denotes observed outcomes.

Among infinitely many interpolating solutions, gradient-based optimization procedures frequently converge toward the minimum-norm solution

$$\theta^* = \arg \min_{\theta} \{\|\theta\|_2 : X\theta = y\}.$$

The minimum-norm solution represents the smoothest or least globally distorted interpolating configuration available within parameter space.

This observation is critical because it demonstrates that smoothness need not emerge from representational scarcity. Smoothness may instead emerge from optimization dynamics operating within highly abundant representational spaces.

The implications are profound.

Under the classical compression paradigm, abstraction is achieved by discarding experiential detail.

Under the excess-capacity framework, abstraction may emerge through dense interpolation itself, provided the representational manifold possesses sufficient flexibility to localize irregularity while preserving global continuity.

The operational consequences of this distinction become most visible when examining the transition between constrained, sufficient, and excess representational regimes. We therefore turn next to a formal analysis of the three-regime structure governing learning behavior.

5 The Three Regimes of Representational Capacity

The central claim of the representational capacity framework is that cognition cannot be adequately understood through static architectural descriptions alone. Cognitive behavior emerges from the dynamic relationship between representational flexibility and experiential complexity. The relevant object of analysis is therefore not the isolated learner, but the evolving geometry between the learner and its experiential environment.

To formalize this relationship, we define the operational capacity ratio

$$\rho = \frac{C}{E},$$

where C denotes effective representational flexibility and E denotes experiential complexity.

The quantity ρ determines the regime in which the learner operates.

This formulation shifts the analysis of intelligence away from absolute size and toward relational structure. No system is intrinsically “large,” “small,” “simple,” or “complex” independent of the experiential manifold it must navigate. A highly expressive system may nevertheless operate under representational scarcity when confronted with sufficiently heterogeneous environments, while comparatively small systems may exhibit representational abundance within restricted domains.

The framework therefore proposes that learning regimes emerge relationally from the geometry between representational flexibility and experiential variability.

5.1 Constrained Capacity Regimes

A constrained regime arises when

$$\rho < 1.$$

In this regime, experiential variability exceeds the representational flexibility available to the learner. The system lacks sufficient degrees of freedom to preserve the detailed structure of its experiences and is therefore forced into compression.

The defining feature of constrained learning is representational sacrifice.

Distinct experiences must be projected onto overlapping internal representations because the learner lacks the dimensional freedom required to preserve fine-grained separation. Local distinctions collapse into broader equivalence classes. Variance is discarded in favor of invariance.

Under constrained capacity, successful generalization emerges through restriction biases.

Restriction biases are structural limitations imposed prior to learning. They constrain the space of admissible representations and thereby force the learner toward compressed abstractions. In polynomial regression, this corresponds to low-order models incapable of representing arbitrary local fluctuations. In

biological cognition, restriction biases may emerge through attentional bottlenecks, cortical compression, metabolic limitations, working-memory constraints, or sparse coding pressures.

Formally, constrained systems seek approximate solutions to

$$f_{\theta}(x_i) \approx y_i,$$

because exact interpolation is impossible under limited representational freedom.

The resulting learned representations exhibit low variance but high bias.

The system generalizes successfully precisely because it cannot memorize every local irregularity. Noise is ignored by necessity rather than by choice.

This produces the classical compression-centered conception of intelligence.

Within this paradigm, abstraction is achieved by discarding experiential detail. Generalization requires forgetting.

Historically, most of cognitive science implicitly operated within this regime assumption. Theories of bounded rationality, prototype abstraction, heuristic simplification, lossy memory consolidation, and statistical regularization all presuppose that cognitive systems lack sufficient representational resources to preserve the detailed structure of experience.

The constrained regime therefore became psychologically normative.

However, this normativity depended upon an unexamined assumption: that successful generalization and detailed memorization are fundamentally incompatible.

The excess-capacity framework rejects this assumption.

5.2 Sufficient Capacity Regimes

A sufficient regime arises near

$$\rho \approx 1.$$

Here, representational flexibility approximately matches experiential complexity. The learner possesses just enough degrees of freedom to interpolate its experiences perfectly, but no substantial excess freedom remains beyond the interpolation constraints themselves.

This regime corresponds to the interpolation threshold.

The sufficient regime is geometrically unstable because local accommodations propagate globally throughout the representational manifold. The learner lacks the surplus dimensionality required to isolate irregularities within localized regions.

Consequently, small perturbations in the experiential distribution produce disproportionately large distortions in the learned representation.

Formally, the system satisfies

$$f_{\theta}(x_i) = y_i$$

for all observed experiences, yet neighboring predictions exhibit extreme sensitivity to perturbation.

This regime corresponds to the classical phenomenon of catastrophic overfitting.

The instability emerges because the interpolation constraints nearly exhaust the representational degrees of freedom available to the learner. Every local correction imposes global consequences.

The sufficient regime therefore occupies a singular position within the geometry of learning. It represents the point at which memorization becomes possible but remains globally brittle.

This observation resolves a long-standing confusion within classical statistical learning theory.

Traditional analyses implicitly conflated interpolation itself with instability. However, interpolation is not inherently pathological. Instability emerges specifically when interpolation occurs under conditions of insufficient surplus representational freedom.

The sufficient regime therefore represents not the inevitability of overfitting, but the pathological boundary between constrained and excess regimes.

This distinction becomes crucial when interpreting double descent phenomena.

The classical U-shaped bias-variance curve correctly characterizes behavior within constrained and near-sufficient regimes. Generalization initially improves as representational flexibility increases, then deteriorates as the learner approaches the interpolation threshold.

However, this curve ceases to describe the system once sufficient capacity is

surpassed.

Beyond the interpolation threshold, the geometry changes fundamentally.

5.3 Excess Capacity Regimes

An excess regime arises when

$$\rho > 1.$$

In this regime, representational flexibility substantially exceeds experiential complexity. The learner possesses more degrees of freedom than are strictly required to interpolate the observed distribution.

This surplus flexibility transforms the geometry of learning.

Because the interpolation constraints no longer exhaust the representational manifold, local irregularities may be accommodated without globally destabilizing representation. The system gains the ability to sequester variance into confined regions while preserving smooth large-scale organization.

The central mechanism of excess-capacity learning is therefore the decoupling of memorization from instability.

To formalize this intuition, consider the decomposition

$$f_{\theta}(x) = g(x) + s(x),$$

where $g(x)$ denotes the globally smooth manifold structure and $s(x)$ denotes localized accommodation structures.

The function $g(x)$ captures broad regularities and stable predictive continuity across experiential space.

The function $s(x)$ captures highly localized deviations required to interpolate particular experiences.

We define a spike as a localized representational accommodation structure whose influence decays rapidly outside a restricted neighborhood of experiential space.

Mathematically, a spike centered near x_i satisfies

$$|s(x)| \rightarrow 0$$

as

$$\|x - x_i\| \rightarrow \infty.$$

The importance of this localization cannot be overstated.

In constrained and sufficient regimes, local irregularities propagate globally because the representational manifold lacks sufficient freedom to isolate them. In excess regimes, however, the learner may preserve detailed episodic structure while maintaining smooth global organization.

This explains how systems may simultaneously memorize and generalize.

Generalization no longer emerges through forced compression. Instead, it emerges through dense interpolation on sufficiently flexible manifolds.

This inversion constitutes the deepest theoretical shift introduced by the framework.

Under the traditional compression paradigm,

memorization \rightarrow failure,

while

abstraction \rightarrow success.

Under excess-capacity learning,

dense interpolation \rightarrow smooth continuity,

and abstraction emerges as a consequence of geometric organization within high-dimensional representational spaces.

The learner does not generalize despite interpolation.

The learner generalizes through interpolation.

5.4 Representational Flexibility and High-Dimensional Separation

The phrase “representations richer than experience” requires careful formal clarification.

Excess-capacity systems do not create information ex nihilo. Information content remains bounded by the entropy of the experiential distribution itself.

What changes is the dimensional organization of representation.

To distinguish these concepts formally, we separate:

$$I(X)$$

as informational entropy,

$$D$$

as representational dimensionality,

and

$$\mathcal{F}$$

as representational flexibility.

A learner may preserve fixed informational content while dramatically increasing representational dimensionality and flexibility.

The result is geometric separation.

Experiences that would collide or interfere within low-dimensional compressed manifolds may remain distinct within higher-dimensional representational spaces.

The role of excess dimensionality is therefore not informational creation but interference avoidance.

This principle aligns with many known properties of overparameterized systems. Kernel methods, transformer architectures, sparse embedding spaces, and high-dimensional associative memory systems all exploit dimensional expansion to preserve separability between otherwise interfering experiential traces.

Biological cognition may operate similarly.

Rather than compressing experience into minimalist summaries whenever possible, neural systems may frequently exploit cortical abundance to preserve local distinctions while maintaining stable global structure.

5.5 The Emergence of Smoothness

A central theoretical question now arises.

Why should excess-capacity systems converge toward smooth solutions at all?

Why do they not simply generate arbitrary chaotic interpolations?

The answer lies in the geometry of high-dimensional optimization landscapes.

In highly overparameterized systems, flat minima occupy vastly larger volumes of parameter space than sharp minima. Optimization procedures based on stochastic search therefore exhibit an emergent probabilistic preference for smooth configurations.

Formally, let

$$\Theta$$

denote parameter space and let

$$\mathcal{M}_{\text{flat}} \subset \Theta$$

represent regions corresponding to low-curvature solutions.

Similarly, let

$$\mathcal{M}_{\text{sharp}} \subset \Theta$$

represent highly localized high-curvature minima.

Then typically,

$$\text{Vol}(\mathcal{M}_{\text{flat}}) \gg \text{Vol}(\mathcal{M}_{\text{sharp}}).$$

Because stochastic optimization trajectories sample parameter neighborhoods probabilistically, they are exponentially more likely to converge toward flat configurations.

Smoothness therefore emerges not merely as an aesthetic preference but as a statistical consequence of high-dimensional geometry.

This principle explains why gradient descent, stochastic gradient descent, and biological plasticity mechanisms frequently converge toward stable generalizing solutions despite massive overparameterization.

The learner is guided toward smooth manifolds because smooth manifolds dominate the accessible geometry of high-dimensional representational space.

5.6 Regime Transitions and Dynamic Cognition

A crucial consequence of the framework is that learning regimes are dynamic rather than fixed.

As experiential complexity increases while representational flexibility remains approximately stable, the operational capacity ratio changes over time:

$$\rho(t) = \frac{C}{E(t)}.$$

This produces regime transitions.

A learner initially operating under excess capacity may gradually transition toward sufficient or constrained regimes as experiential accumulation overwhelms available representational flexibility.

This dynamic explains several otherwise puzzling developmental and cognitive phenomena.

Children may initially memorize irregular linguistic forms because early experiential complexity remains low relative to representational flexibility. As vocabulary grows rapidly, however, experiential accumulation may temporarily exceed available representational abundance, forcing the learner into constrained rule-based compression.

The child therefore transitions from:

went → goed → went.

The apparent regression reflects not irrationality but regime transition.

Similarly, perceptual deterioration after extensive training may emerge because accumulated experiences consume representational surplus previously available for smoothness stabilization.

Learning itself alters the geometry of cognition.

The mind is therefore not best modeled as a fixed architecture of static modules. It is more accurately understood as a dynamically evolving representational state-space navigating shifting capacity regimes across time, task, and environment.

6 Double Descent and Benign Overfitting

The most direct empirical challenge to the historical compression paradigm emerges from the phenomenon known as double descent. Classical statistical learning theory predicted that increasing representational complexity beyond a certain threshold would inevitably produce catastrophic overfitting. Generalization error

was expected to follow a U-shaped curve. Initially, increasing model flexibility reduces bias by allowing the learner to capture increasingly rich structure. However, once representational flexibility becomes sufficiently large to interpolate the training data exactly, variance explodes and generalization collapses.

For decades, this relationship appeared foundational.

The classical bias–variance tradeoff became deeply integrated into psychology, statistics, neuroscience, and machine learning. Simplicity was not merely preferred; it was presumed mathematically necessary for robust inference.

Yet modern overparameterized systems repeatedly violate this prediction.

Deep neural networks containing vastly more parameters than training examples frequently interpolate noisy datasets perfectly while continuing to generalize effectively. Rather than diverging catastrophically after interpolation, test error often decreases again once representational flexibility substantially exceeds exponential complexity.

This produces the characteristic double descent curve.

To understand why this phenomenon is theoretically transformative, we first examine the geometry underlying the classical U-curve.

6.1 The Classical Bias–Variance Paradigm

Consider a learner attempting to approximate an unknown latent function

$$f^*(x)$$

from noisy observations

$$y_i = f^*(x_i) + \epsilon_i,$$

where

$$\epsilon_i$$

denotes stochastic noise.

Suppose the learner constructs an approximation

$$f_\theta(x)$$

parameterized by

$$\theta.$$

The expected prediction error may traditionally be decomposed into

$$\mathbb{E} [(y - f_{\theta}(x))^2] = \text{Bias}^2 + \text{Variance} + \sigma^2,$$

where

$$\sigma^2$$

denotes irreducible noise.

Within this framework, constrained models exhibit high bias because they cannot represent sufficiently rich structure. Increasing representational flexibility initially improves performance by reducing approximation error.

However, classical theory predicts that once the learner acquires enough flexibility to interpolate the training data exactly, variance increases catastrophically.

The learner begins fitting noise rather than structure.

Generalization deteriorates.

This interpretation treats interpolation itself as the source of instability.

The excess-capacity framework argues that this conclusion is incomplete.

6.2 The Interpolation Threshold

The crucial distinction introduced by modern learning theory is that interpolation and instability are not synonymous.

Instability emerges specifically at the interpolation threshold.

Let the learner possess effective representational dimensionality

$$d.$$

Suppose the dataset contains

$$n$$

independent interpolation constraints.

The interpolation threshold approximately occurs when

$$d \approx n.$$

At this point, the learner possesses just enough flexibility to satisfy every

constraint exactly but lacks substantial surplus freedom beyond those constraints.

The geometry becomes maximally brittle.

To illustrate this formally, consider linear interpolation under the constraint

$$X\theta = y,$$

where

$$X \in \mathbb{R}^{n \times d}$$

denotes the design matrix.

When

$$d < n,$$

exact interpolation is generally impossible. The learner must compress.

When

$$d \approx n,$$

the solution space narrows dramatically. Small perturbations in the observations produce disproportionately large changes in the interpolating solution.

The resulting interpolants exhibit extreme curvature and instability.

This is the classical overfitting regime.

Importantly, however, this instability is not caused by interpolation itself. It is caused by attempting interpolation without sufficient surplus representational freedom.

The distinction becomes visible once

$$d \gg n.$$

6.3 The Second Descent

In highly overparameterized systems, the geometry changes fundamentally.

When representational flexibility substantially exceeds the number of interpolation constraints, the learner possesses infinitely many exact interpolating solutions.

The optimization problem becomes underdetermined.

Instead of being forced into globally unstable distortions, the learner may

localize irregular accommodations within restricted regions of representational space while preserving smooth large-scale structure elsewhere.

The generalization error therefore decreases again.

This phenomenon constitutes the second descent.

The crucial insight is that overparameterization transforms interpolation from a globally constrained problem into a geometrically flexible one.

Classical theory assumed that memorization necessarily destroys abstraction because it studied systems operating near the interpolation threshold.

Double descent demonstrates that the interpolation threshold is not the final regime of learning geometry. It is merely the unstable boundary separating constrained and excess-capacity behavior.

6.4 Polynomial Interpolation and Geometric Intuition

The intuition underlying double descent becomes particularly clear when examining polynomial interpolation.

Suppose the learner attempts to fit noisy observations sampled from an underlying smooth latent function.

A degree-one polynomial lacks sufficient flexibility to interpolate the observations exactly. The resulting approximation is smooth but biased.

A polynomial near degree

$$n - 1$$

possesses just enough flexibility to pass through every observation. However, satisfying all interpolation constraints requires extreme oscillatory deformation throughout the domain.

The interpolating curve becomes globally unstable.

This corresponds to the sufficient-capacity regime.

Now consider a polynomial or function class with vastly greater flexibility than required for interpolation.

Instead of globally oscillating to satisfy every local irregularity, the learner may generate highly localized accommodation structures concentrated near the noisy observations themselves.

The result is qualitatively different.

The learner still interpolates perfectly:

$$f_{\theta}(x_i) = y_i,$$

yet the global manifold remains smooth almost everywhere else.

Noise becomes localized rather than globally propagated.

The classical interpretation of overfitting therefore requires revision.

Overfitting is not intrinsically catastrophic.

Overfitting becomes catastrophic specifically when the learner lacks sufficient representational surplus to isolate irregularity geometrically.

6.5 Benign Overfitting

This reinterpretation gives rise to the concept of benign overfitting.

Benign overfitting refers to situations in which a learner perfectly interpolates noisy observations while maintaining low generalization error.

At first glance, this appears paradoxical. Classical statistical learning theory treated interpolation and robustness as fundamentally incompatible.

However, within the excess-capacity framework, the paradox dissolves naturally.

Suppose the learner decomposes its representation into globally smooth structure and localized residual accommodations:

$$f_{\theta}(x) = g(x) + s(x),$$

where

$$g(x)$$

captures smooth manifold continuity and

$$s(x)$$

captures localized deviations required for exact interpolation.

In excess-capacity systems, the residual structures remain geometrically confined:

$$|s(x)| \rightarrow 0$$

outside restricted neighborhoods surrounding noisy observations.

The learner therefore memorizes local irregularities without allowing those irregularities to globally distort predictive structure.

Memorization and generalization become decoupled.

This decoupling represents one of the central theoretical claims of the framework.

The historical compression paradigm assumed:

$$\text{memorization} \Rightarrow \text{loss of generalization.}$$

Excess-capacity learning instead permits:

$$\text{memorization} \wedge \text{generalization.}$$

The compatibility emerges not from ignoring variance, but from localizing it.

6.6 Minimum-Norm Interpolation

A central question now arises.

If infinitely many interpolating solutions exist in excess-capacity systems, why do learners preferentially converge toward smooth solutions rather than arbitrary chaotic interpolants?

The answer lies in implicit regularization.

Modern optimization procedures frequently converge toward minimum-norm interpolating solutions even in the absence of explicit regularization constraints.

Formally, consider the interpolation problem

$$X\theta = y.$$

Among infinitely many admissible interpolants, optimization procedures such as gradient descent often converge toward

$$\theta^* = \arg \min_{\theta} \{\|\theta\|_2 : X\theta = y\}.$$

This solution minimizes global parameter magnitude while satisfying all inter-

polation constraints exactly.

The resulting interpolant exhibits reduced global curvature and greater manifold smoothness.

The importance of this observation cannot be overstated.

Smoothness does not emerge because the learner lacks flexibility.

Smoothness emerges because optimization geometry favors stable interpolating manifolds among the enormous set of admissible exact solutions.

This distinction fundamentally separates excess-capacity learning from classical compression-centered cognition.

Constrained systems generalize because they are forced to simplify.

Excess systems generalize because high-dimensional optimization landscapes probabilistically favor smooth interpolation geometries.

6.7 Volume Effects and High-Dimensional Geometry

The emergence of smooth solutions in overparameterized systems may be understood through volume effects in high-dimensional spaces.

Let

$$\Theta$$

denote parameter space.

Suppose

$$\mathcal{M}_{\text{flat}} \subset \Theta$$

corresponds to low-curvature solutions and

$$\mathcal{M}_{\text{sharp}} \subset \Theta$$

corresponds to highly localized sharp minima.

In high-dimensional systems,

$$\text{Vol}(\mathcal{M}_{\text{flat}}) \gg \text{Vol}(\mathcal{M}_{\text{sharp}}).$$

Flat minima occupy exponentially larger regions of parameter space.

Consequently, stochastic optimization procedures are overwhelmingly more likely to converge toward smooth solutions.

This principle provides the deep geometric explanation for benign overfitting.

Smoothness is not merely an imposed preference.

Smoothness emerges probabilistically because stable low-curvature interpolants dominate the geometry of high-dimensional representational spaces.

6.8 The Philosophical Inversion

The implications of double descent extend far beyond machine learning.

The phenomenon forces a profound inversion in the ontology of intelligence itself.

Under the historical compression paradigm, abstraction emerged through simplification. Intelligence meant reducing the world to compact summaries while discarding noisy experiential detail.

Double descent reveals another possibility.

Abstraction may emerge through dense interpolation itself.

Highly expressive systems may preserve enormous quantities of local detail while still converging toward globally smooth predictive organization.

The learner does not succeed by eliminating complexity.

The learner succeeds because excess representational freedom allows complexity to become geometrically organized.

This inversion transforms the meaning of intelligence.

The mind is no longer understood primarily as a machine built to squeeze the world into compressed symbolic summaries.

Instead, cognition becomes the problem of maintaining stable global continuity while preserving high-fidelity local accommodation within richly overparameterized representational spaces.

The resulting picture is fundamentally different from the economy-of-thought tradition that dominated the twentieth century.

Complexity is no longer the enemy of generalization.

Under conditions of representational abundance, complexity becomes its substrate.

7 Smooth Interpolation and Soft Inductive Bias

The previous section established that highly overparameterized systems can interpolate noisy observations while nevertheless preserving strong generalization performance. This observation destabilizes the classical assumption that interpolation and robustness are fundamentally incompatible. However, the phenomenon raises a deeper theoretical question.

Why should excess-capacity systems converge toward smooth, stable, and generalizable interpolants at all?

If infinitely many exact solutions exist, most of them should intuitively appear pathological. A learner with unrestricted flexibility could in principle memorize observations through arbitrarily jagged and unstable representations. Yet empirical systems repeatedly converge toward solutions exhibiting substantial smoothness and predictive continuity.

The answer lies in the geometry of optimization itself.

The excess-capacity framework argues that smoothness emerges not from representational scarcity, but from soft inductive biases generated by high-dimensional optimization dynamics.

This distinction is foundational.

Classical compression-centered theories derive smoothness from incapacity. The learner generalizes because it lacks sufficient representational freedom to preserve arbitrary local irregularities.

Excess-capacity systems operate differently. They possess many more representational degrees of freedom than are strictly required for interpolation. Generalization therefore cannot emerge from restriction alone. Instead, it emerges from probabilistic preferences governing the geometry of admissible interpolating solutions.

7.1 Restriction Biases and Soft Biases

To clarify the distinction formally, we define two fundamentally different mechanisms of generalization.

A restriction bias is a hard architectural limitation that constrains the space of possible representations available to the learner.

Examples include low-dimensional parameterizations, sparse feature con-

straints, low-order polynomials, compressed latent bottlenecks, or biologically imposed resource limits.

Restriction biases force simplification because the learner literally cannot represent arbitrary complexity.

Soft inductive biases, by contrast, emerge from optimization geometry within highly expressive systems.

The learner possesses many admissible interpolating solutions:

$$\mathcal{S} = \{\theta : X\theta = y\}.$$

The optimization process then selects among these admissible solutions according to probabilistic, geometric, or dynamical preferences.

The critical insight is that smoothness may emerge from selection among many exact solutions rather than from an inability to represent alternatives.

This transforms the ontology of abstraction itself.

7.2 Interpolation Geometry

Suppose a learner seeks an interpolating function

$$f_{\theta}(x)$$

satisfying

$$f_{\theta}(x_i) = y_i$$

for all observed experiences.

When the system operates in an excess-capacity regime,

$$\rho > 1,$$

the interpolation constraints no longer uniquely determine the solution.

The solution set becomes highly degenerate.

The learner therefore navigates an enormous family of exact interpolants.

The geometry of this solution space is crucial.

Let

$$\Theta$$

denote parameter space and let

$$\mathcal{S} \subset \Theta$$

denote the manifold of exact interpolating solutions.

Within constrained or sufficient regimes, the interpolation manifold remains narrow or nonexistent. The learner is forced toward globally unstable solutions because representational freedom is exhausted by the interpolation constraints themselves.

In excess-capacity regimes, however,

$$\dim(\mathcal{S}) \gg 0.$$

The solution manifold becomes extremely high-dimensional.

The learner is no longer solving a single interpolation problem.

Instead, it is selecting among vast families of admissible representational geometries.

This distinction is the key to benign overfitting.

7.3 Minimum-Norm Dynamics

One of the most important empirical findings in modern learning theory is that optimization procedures frequently converge toward low-complexity interpolants even when no explicit regularization term is imposed.

Consider again the interpolation problem

$$X\theta = y.$$

Among infinitely many admissible solutions, gradient-based optimization often converges toward

$$\theta^* = \arg \min_{\theta} \{\|\theta\|_2 : X\theta = y\}.$$

This minimum-norm interpolant minimizes global parameter magnitude while preserving exact interpolation.

The resulting geometry exhibits reduced curvature, increased smoothness, and improved stability.

Importantly, this preference is not explicitly programmed into the learner. It emerges naturally from optimization dynamics.

The learner therefore behaves as though it “prefers” smoothness, even though smoothness is not externally imposed.

This phenomenon is known as implicit regularization.

Implicit regularization represents one of the most significant conceptual developments in modern machine learning because it demonstrates that optimization procedures themselves may induce strong structural preferences.

The learner does not generalize because complexity is impossible.

The learner generalizes because optimization trajectories disproportionately converge toward smooth regions of representational space.

7.4 Flat Minima and Volume Effects

The emergence of smooth interpolants may be understood geometrically through volume effects in high-dimensional parameter spaces.

Suppose parameter space

$$\Theta$$

contains two broad classes of solutions.

The first class consists of sharp minima:

$$\mathcal{M}_{\text{sharp}},$$

corresponding to highly localized and unstable solutions.

The second class consists of flat minima:

$$\mathcal{M}_{\text{flat}},$$

corresponding to low-curvature stable interpolants.

Empirically and theoretically,

$$\text{Vol}(\mathcal{M}_{\text{flat}}) \gg \text{Vol}(\mathcal{M}_{\text{sharp}}).$$

Flat minima occupy vastly larger neighborhoods of parameter space.

This asymmetry has profound consequences.

Stochastic optimization procedures such as stochastic gradient descent sample parameter trajectories probabilistically. Because flat regions occupy much larger volumes, optimization dynamics are overwhelmingly more likely to converge toward smooth configurations.

The preference for smoothness therefore emerges statistically from the geometry of high-dimensional spaces themselves.

Smoothness is not an externally imposed principle.

It is an attractor property of abundant representational geometry.

7.5 Localized Residual Accommodation

The interaction between global smoothness and local memorization now becomes formally intelligible.

Suppose the learner decomposes its representation into

$$f_{\theta}(x) = g(x) + s(x),$$

where

$$g(x)$$

denotes the globally smooth manifold structure and

$$s(x)$$

denotes localized accommodation structures.

The function

$$g(x)$$

captures broad predictive continuity.

The function

$$s(x)$$

captures highly localized corrections required to satisfy specific interpolation constraints.

We define a localized residual accommodation structure as a representational deviation whose influence decays rapidly outside a restricted neighborhood of experiential space.

Formally,

$$|s(x)| \rightarrow 0$$

as

$$\|x - x_i\| \rightarrow \infty.$$

These localized structures correspond to what earlier sections termed “spikes.”

The critical feature of spikes is that they preserve local detail without globally destabilizing the manifold.

This localization explains how memorization and generalization may coexist.

The learner preserves highly specific experiential traces while maintaining smooth predictive organization almost everywhere else.

The result is not a tradeoff between rules and exceptions.

Instead, rules emerge as large-scale geometric continuity within manifolds capable of accommodating localized exceptions.

7.6 Representational Expansion and Interference Avoidance

The smooth interpolation framework also clarifies the meaning of representational abundance.

A high-capacity system does not create information beyond the entropy of the experiential distribution itself.

Instead, excess representational dimensionality allows the learner to reduce interference between experiential traces.

Suppose two experiences

$$x_i \quad \text{and} \quad x_j$$

occupy nearby regions within low-dimensional representational space.

In constrained systems, these experiences may collide representationally, forcing the learner to compress distinctions between them.

Increasing representational dimensionality allows the learner to separate these experiences geometrically.

The result is not informational creation but spatial separation.

Representational abundance therefore reduces the pressure to compress.

This principle aligns naturally with kernel methods, sparse distributed mem-

ory, transformer attention structures, cortical expansion, and high-dimensional embedding systems.

The learner preserves detail because abundant representational geometry permits experiential traces to remain separated rather than collapsed into compressed summaries.

7.7 The Emergence of Abstraction

One of the deepest consequences of the framework is that abstraction itself must now be reinterpreted.

Under classical compression-centered cognition, abstraction emerges by eliminating variance.

The learner discovers invariant structure precisely by discarding local irregularity.

Under excess-capacity learning, however, abstraction emerges differently.

Abstraction becomes the global continuity structure of densely interpolated manifolds.

The learner preserves enormous quantities of local experiential detail while still exhibiting large-scale predictive smoothness.

The global manifold itself becomes the abstraction.

This inversion fundamentally alters the relationship between memorization and intelligence.

The classical view assumed:

memorization \Rightarrow loss of abstraction.

The excess-capacity framework instead permits:

dense interpolation \Rightarrow emergent smoothness.

Generalization therefore becomes a geometric property of representational organization rather than merely the consequence of forced simplification.

7.8 Biological Implications

The emergence of smooth interpolation has significant implications for biological cognition.

Neural systems exhibit extensive evidence of representational expansion. Sensory pathways frequently project lower-dimensional inputs into substantially higher-dimensional cortical spaces. Sparse coding, distributed representation, recurrent connectivity, and population coding all increase representational flexibility relative to raw sensory input.

These properties are difficult to reconcile with purely compression-centered theories.

Within the excess-capacity framework, however, such expansion becomes computationally intelligible.

Biological cognition may preserve abundant representational flexibility precisely because abundant geometry supports robust interpolation while minimizing catastrophic interference.

The brain may therefore operate not merely as a compression engine, but as a high-dimensional manifold management system balancing local accommodation against global continuity.

This reinterpretation also weakens the necessity of strong modular separation between episodic and statistical learning systems.

If sufficiently expressive representational manifolds can simultaneously preserve local detail and global continuity, then dual-process phenomena may emerge from different operating regimes within unified representational substrates rather than from fundamentally distinct learning architectures.

The implications of this claim become particularly important when examining human empirical signatures of excess-capacity learning. We therefore turn next to the behavioral, developmental, and cognitive phenomena suggesting that biological learners frequently operate under conditions of representational abundance.

8 Human Empirical Signatures of Excess Capacity

The theoretical framework developed thus far proposes that biological cognition often operates within regimes of representational abundance rather than strict compression. If this claim is correct, then the empirical signatures of human

learning should differ substantially from those predicted by classical parsimony-centered theories. Human cognition should exhibit not merely abstraction and compression, but persistent retention of detail, sensitivity to local structure, dependence on experiential geometry, and forms of generalization emerging through dense interpolation rather than reductive filtering.

This section examines the empirical phenomena most consistent with excess-capacity learning and develops a unified interpretation of these observations within the regime framework introduced earlier.

8.1 The Classical Compression Prediction

Under traditional simplicity-centered theories, successful cognition requires aggressive elimination of variance.

Suppose a learner encounters experiences

$$\mathcal{E} = \{(x_i, y_i)\}_{i=1}^n.$$

A constrained learner possessing insufficient representational flexibility

$$C$$

relative to experiential complexity

$$E$$

must minimize representational burden by extracting only invariant statistical structure.

Formally, constrained systems attempt to minimize:

$$I(R(X); X),$$

where

$$R(X)$$

denotes the compressed representation of experience.

The compression-centered prediction therefore implies several behavioral properties.

First, memory for arbitrary detail should decay rapidly.

Second, superficial variance should be ignored whenever abstract rules are

available.

Third, increasing exposure should monotonically improve abstraction through progressive elimination of idiosyncratic noise.

Fourth, memorization and generalization should compete directly for representational resources.

The empirical literature repeatedly violates these predictions.

8.2 Massive Recognition Fidelity

One of the strongest empirical signatures of excess-capacity learning is the enormous recognition fidelity exhibited by biological cognition.

Brady et al. demonstrated that participants exposed to approximately

2500

visual objects for only a few seconds each later recognized them with remarkable precision, even when distractors belonged to highly similar semantic categories.

Classical compression models struggle to explain this result.

If cognition generalized primarily through aggressive variance reduction, then fine-grained object-specific detail should have been discarded rapidly as informationally irrelevant.

Instead, human learners preserved large quantities of arbitrary experiential structure.

The excess-capacity interpretation explains this naturally.

Suppose experiential traces are embedded into a high-dimensional representational manifold:

$$\phi : X \rightarrow \mathbb{R}^d,$$

with

$$d \gg n.$$

The learner is no longer forced to collapse nearby experiences into compressed summaries because representational dimensionality permits substantial geometric separation between traces.

Distinct experiences occupy separable regions of representational space:

$$\|\phi(x_i) - \phi(x_j)\| \gg 0.$$

Memory therefore persists not because compression has become unnecessary, but because interference pressure has been reduced through representational abundance.

This predicts a key property of excess-capacity systems:

generalization and memorization become partially decoupled.

The learner can preserve local detail while still extracting globally smooth predictive structure.

8.3 Superficial Detail Sensitivity

Further evidence against strict compression appears in human reliance upon apparently unnecessary detail.

Lupyan demonstrated that human category judgments frequently incorporate superficial or statistically irrelevant features even when simpler categorical rules are explicitly available.

Participants identifying triangles, for example, often relied upon side length proportions, orientation cues, or prototypical appearance despite possessing the abstract rule:

$$\text{triangle} \iff \text{three-sided polygon}.$$

Under constrained-capacity theories, such behavior appears inefficient.

The learner should discard unnecessary variance and preserve only minimal sufficient structure.

Yet humans repeatedly preserve local representational richness beyond what strict abstraction would require.

The excess-capacity framework predicts precisely this phenomenon.

Dense interpolation systems preserve local accommodation structures because representational pressure to collapse distinctions remains low.

Rather than compressing all exemplars into minimal prototypes, the learner constructs globally smooth manifolds populated by highly differentiated local traces.

Consequently, categorical reasoning becomes simultaneously abstract and exemplar-sensitive.

This duality previously motivated strong modular theories separating rule-based and exemplar-based cognition into distinct systems.

The regime framework instead predicts that both behaviors may emerge naturally within unified excess-capacity substrates.

8.4 Prototype Effects and Emergent Smoothness

Although excess-capacity systems preserve local detail, they nevertheless exhibit powerful prototype effects.

Posner and Keele famously demonstrated that participants classified unseen prototype stimuli more efficiently than previously encountered exemplars.

At first glance, this appears to support classical compression theories.

However, within the regime framework, prototype formation emerges naturally through manifold geometry rather than through variance elimination.

Suppose the learner constructs a smooth interpolating manifold:

$$\mathcal{M} \subset \mathbb{R}^d.$$

Individual exemplars correspond to localized regions on the manifold:

$$x_i \in \mathcal{M}.$$

The prototype then corresponds approximately to a low-curvature central region minimizing aggregate representational distance:

$$x^* = \arg \min_x \sum_i \|x - x_i\|^2.$$

Importantly, this prototype emerges without requiring destruction of exemplar-specific detail.

The learner preserves both:

global smoothness

and

local accommodation.

Prototype effects therefore do not imply aggressive compression.

They may instead reflect geometric averaging within high-dimensional interpolation manifolds.

This reinterpretation constitutes one of the major conceptual shifts of excess-capacity learning.

Abstraction becomes an emergent geometric property of densely represented experience rather than a consequence of forced forgetting.

8.5 Perceptual Deterioration and Regime Transition

One of the most counterintuitive predictions of the regime framework is that additional experience may sometimes impair generalization.

This prediction sharply contrasts with classical learning theories, which assume monotonic improvement with increased exposure.

Censor and Sagi observed precisely such perceptual deterioration effects within visual discrimination tasks.

Participants receiving excessive training eventually exhibited reduced transfer performance despite continued exposure.

Under classical models, this result appears paradoxical.

The excess-capacity framework explains the phenomenon naturally through dynamic regime transition.

Recall that capacity is relational:

$$\rho = \frac{C}{E}.$$

Suppose representational flexibility

C

remains approximately fixed while experiential complexity

E

continues increasing.

Then:

$$\frac{d\rho}{dE} < 0.$$

As additional experiences accumulate, the learner progressively loses excess representational freedom.

The system therefore approaches the interpolation threshold:

$$\rho \approx 1.$$

Near this threshold, local accommodation structures begin interfering globally because insufficient excess geometry remains available to isolate them cleanly.

The manifold becomes increasingly brittle.

Smoothness deteriorates.

Generalization weakens.

Perceptual deterioration thus becomes a direct signature of regime transition rather than evidence against learning itself.

This prediction is especially important because it demonstrates that excess-capacity learning is not triumphalist.

Representational abundance confers advantages only under specific relational conditions.

8.6 False Recognition and Representational Blurring

Koutstaal and Schacter demonstrated that humans frequently exhibit false recognition for semantically related but previously unseen category members.

Classical theories often interpret this as evidence for prototype-based abstraction replacing episodic detail.

The regime framework offers a more nuanced interpretation.

Suppose local experiential traces are initially represented through highly differentiated accommodation structures:

$$s_i(x).$$

When excess capacity is abundant:

$$\rho \gg 1,$$

the learner preserves substantial geometric separation between neighboring exemplars.

As experiential load increases relative to representational flexibility, however, local trace separation becomes increasingly difficult to maintain.

The manifold gradually smooths over previously distinct regions:

$$s_i(x) \approx s_j(x).$$

The learner therefore begins confusing nearby category members because representational neighborhoods overlap more strongly.

False recognition emerges not because exemplars disappeared entirely, but because representational abundance has weakened relative to experiential density.

This interpretation predicts that false recognition rates should vary dynamically with relative capacity conditions rather than reflecting fixed architectural separation between episodic and semantic systems.

8.7 Noise Sensitivity and Adversarial Fragility

Excess-capacity systems exhibit another distinctive property:

high sensitivity to local noise.

Constrained learners frequently ignore noisy observations because representational scarcity prevents accommodation of arbitrary deviations.

Excess-capacity learners behave differently.

Because interpolation pressure remains high, the learner attempts to account for nearly every local fluctuation.

Suppose noisy observations satisfy:

$$y_i = f(x_i) + \epsilon_i.$$

The excess-capacity learner constructs:

$$\hat{f}(x_i) = y_i$$

exactly, producing localized accommodation structures:

$$s_i(x).$$

This increases sensitivity to adversarial perturbation, spurious correlations, and local distributional anomalies.

The system remains globally smooth only if local accommodations remain sufficiently isolated.

When environmental noise becomes too dense or highly structured, accommodation structures begin interfering globally, degrading generalization.

This predicts an important boundary condition:

constrained systems may outperform excess systems under highly adversarial or unstable environments.

Compression therefore remains adaptive under specific ecological conditions.

The framework thus becomes a regime theory rather than an ideological rejection of parsimony itself.

8.8 Developmental U-Shaped Learning

One of the most compelling developmental signatures of regime dynamics appears in U-shaped learning phenomena.

Children frequently progress from correct irregular forms:

went,

to overgeneralized forms:

goed,

before eventually returning to correct irregular usage.

Classical dual-process accounts interpret this as competition between memorized exceptions and abstract grammatical rules.

The excess-capacity framework instead interprets this as a dynamic shift in

representational regime.

Early in development, experiential complexity remains low:

$$E \ll C.$$

The learner therefore operates within an excess-capacity regime capable of preserving isolated irregular exemplars.

As vocabulary expands rapidly:

$$E \uparrow,$$

relative capacity decreases:

$$\rho \downarrow.$$

The learner temporarily loses sufficient excess geometry to preserve all isolated irregular traces.

Global grammatical smoothness dominates local accommodation structures.

The system therefore overgeneralizes:

$$\text{go} \mapsto \text{goed}.$$

As representational flexibility later expands through developmental growth and accumulated organization, irregular traces become reintegrated into a more stable interpolation manifold.

The learner returns to:

$$\text{went}.$$

The U-shaped trajectory therefore reflects dynamic movement between representational regimes rather than conflict between separate cognitive systems.

8.9 Attention, Constraint, and Gradient Preservation

The implications of excess-capacity learning extend beyond memory and categorization into attentional dynamics themselves.

Constraint-driven attentional ecologies appear especially consistent with representational abundance frameworks. Systems maintaining sustained engagement frequently preserve dense unresolved structure rather than aggressively compress-

ing experience into completed summaries.

Constraint-first learning practices maintain informational gradients by preserving partially unresolved manifolds rather than terminating representation through closure.

Attention persists because local curvature remains available for further descent.

Under this interpretation, boredom corresponds to flattened representational geometry:

$$\nabla \mathcal{M} \approx 0.$$

Engagement persists when representational manifolds remain sufficiently differentiated to sustain continued interpolation.

This interpretation aligns naturally with broader geometric and constraint-driven theories of cognition, where learning becomes motion through structured representational fields rather than static compression into minimal symbolic summaries.

8.10 Toward a Unified Empirical Interpretation

Taken together, the empirical phenomena reviewed in this section strongly suggest that human cognition frequently operates under conditions incompatible with strict compression-centered learning.

Humans preserve enormous quantities of arbitrary detail.

They remain sensitive to local experiential structure even when abstract rules are available.

They exhibit perceptual deterioration under excessive exposure.

They produce prototype effects without eliminating exemplar specificity.

They demonstrate developmental regime transitions.

They remain vulnerable to local noise and adversarial perturbation.

These properties emerge naturally under excess-capacity learning.

The learner generalizes not by eliminating local detail, but by embedding local detail within globally smooth representational manifolds possessing sufficient dimensional flexibility to preserve both simultaneously.

The implications of this reinterpretation extend far beyond isolated behavioral experiments.

They challenge the necessity of strong modular architectures separating memory from abstraction, rules from exemplars, or statistical learning from episodic retention.

The next section therefore examines how excess-capacity learning reconfigures broader theories of cognition, development, and unified representational architecture.

9 Developmental Dynamics and Regime Transitions

One of the most important consequences of the representational capacity framework is that cognitive behavior is not fixed. Learning systems do not permanently inhabit a single operational regime. Instead, they move dynamically through different regions of representational organization as experiential complexity, environmental structure, developmental maturation, and representational flexibility evolve over time.

This transforms cognition from a static architecture into a trajectory through a representational phase space.

Under classical theories, developmental changes are often interpreted as transitions between modules, strategies, or qualitatively distinct reasoning systems. Rule-based cognition replaces memorization, abstraction supersedes episodic retention, or symbolic reasoning emerges from associative learning through architectural differentiation.

The excess-capacity framework proposes a different interpretation.

Many developmental phenomena may instead emerge from continuous shifts in the relational ratio:

$$\rho = \frac{C}{E},$$

where

$$C$$

represents effective representational flexibility and

$$E$$

represents experiential complexity.

Under this interpretation, developmental transitions correspond not primarily to replacement of mechanisms, but to movement between operational geometries.

9.1 Development as Representational Rebalancing

Suppose a learner begins life with relatively sparse experiential exposure:

$$E_0 \ll C_0.$$

The system therefore initially occupies an excess-capacity regime:

$$\rho \gg 1.$$

At this stage, the learner possesses substantial unused representational flexibility relative to environmental complexity. Local experiential traces may therefore remain highly differentiated.

This predicts several early developmental properties.

First, episodic specificity should remain strong.

Second, arbitrary local distinctions should exert disproportionate influence.

Third, learners should display high memorization fidelity despite relatively weak global abstraction.

Empirical developmental evidence strongly supports these predictions.

Young learners frequently retain arbitrary surface regularities, imitate highly specific contextual details, and preserve local experiential traces with remarkable precision.

However, development rapidly increases experiential complexity:

$$E(t) \uparrow .$$

Vocabulary expands, social environments become more intricate, causal structure becomes denser, and representational demands multiply.

If representational flexibility grows more slowly than experiential complexity, then:

$$\frac{d\rho}{dt} < 0.$$

The learner therefore gradually moves toward constrained or near-interpolation regimes.

Developmental cognition consequently changes not because one cognitive

system replaces another, but because the geometry of representational abundance itself evolves.

9.2 The Geometry of U-Shaped Development

U-shaped developmental trajectories provide one of the clearest signatures of dynamic regime transition.

Children often initially produce correct irregular forms:

went, came, broke,

before later overgeneralizing:

goed, comed, breaked,

and eventually returning to correct irregular usage.

Traditional dual-process theories interpret this phenomenon as competition between memorized exceptions and abstract grammatical rules.

The representational regime framework offers a more unified account.

Early in development, low experiential density permits preservation of isolated irregular accommodation structures:

$$s_i(x).$$

Irregular verbs remain represented as distinct local attractors embedded within sparse experiential geometry.

As linguistic exposure increases rapidly:

$$E \uparrow,$$

relative excess capacity decreases:

$$\rho \downarrow.$$

The learner gradually loses sufficient representational flexibility to preserve every irregular structure independently.

Global manifold smoothness therefore begins dominating local accommoda-

tion:

$$\text{past tense} \approx \text{verb} + \text{ed}.$$

The learner consequently collapses irregular forms into broader grammatical continuity structures.

Overgeneralization emerges naturally.

Eventually, representational organization improves further through cortical maturation, manifold differentiation, and increasing representational efficiency. The learner regains sufficient geometric separation to preserve both:

global grammatical smoothness

and

localized irregular exceptions.

Correct irregular forms then re-emerge.

The U-shaped curve therefore reflects representational phase transition rather than competition between distinct symbolic and associative systems.

9.3 Constraint Emergence Through Experiential Saturation

An especially important implication follows from the dynamic nature of capacity.

Increasing experience may force abstraction even without explicit compression objectives.

Suppose representational flexibility remains approximately bounded:

$$C \approx \text{constant},$$

while experiential density continues increasing:

$$E \rightarrow \infty.$$

Then:

$$\rho \rightarrow 0.$$

Under such conditions, representational abundance disappears entirely.

The learner becomes progressively constrained by experiential saturation.

This predicts that abstraction itself may emerge as an adaptive response to representational crowding.

The learner begins compressing because local differentiation becomes geometrically unsustainable.

This reinterpretation changes the ontology of abstraction.

Under classical theories, abstraction is primary and memorization secondary.

Under the regime framework, memorization may initially dominate. Abstraction emerges gradually as representational density forces manifold simplification.

The child does not begin as an abstract reasoner.

The child begins as a dense interpolator operating within sparse experiential geometry.

9.4 Play, Constraint, and Generative Differentiation

The framework also provides a new interpretation of exploratory play.

Traditional developmental theories often treat play as loosely structured experimentation or intrinsically motivated behavioral variation.

The representational regime framework instead interprets play as active geometric differentiation under conditions of relatively abundant representational freedom.

Young learners possess relatively low experiential density:

E

small, combined with substantial representational flexibility:

C

moderately large.

The resulting excess-capacity regime permits broad exploratory interpolation without immediate pressure toward aggressive compression.

Play therefore becomes a mechanism for constructing representational curvature.

The learner imposes temporary voluntary constraints:

$$\mathcal{C}_1, \mathcal{C}_2, \dots$$

upon otherwise weakly differentiated environments.

Games, symbolic play, imaginative structures, linguistic experimentation, and exploratory recombination all function by selectively tightening local manifolds to generate informational gradients.

This predicts an important developmental principle.

Attention stabilizes when representational manifolds remain partially unresolved.

Play succeeds not because it maximizes stimulation, but because it maintains generative incompleteness.

The learner remains engaged because the representational field continues producing admissible differentiation.

This interpretation aligns naturally with broader constraint-first approaches to cognition and learning, where sustained engagement emerges from structured incompleteness rather than passive consumption.

9.5 The Development of Heuristics

As representational density increases, learners progressively rely upon sparse coordinate systems for navigating increasingly complex manifolds.

Heuristics emerge naturally under these conditions.

Suppose the learner initially represents experiential space using highly differentiated local embeddings:

$$\phi(x_i).$$

As experiential complexity grows, maintaining exhaustive pairwise differentiation becomes increasingly costly.

The learner therefore constructs lower-dimensional coordinate projections:

$$\psi : \mathbb{R}^d \rightarrow \mathbb{R}^k, \quad k \ll d.$$

Heuristics correspond to these compressed coordinate charts.

Importantly, heuristics are not necessarily failures of rationality.

They are adaptive responses to changing regime geometry.

Sparse projections preserve navigationally relevant gradients while discarding dimensions whose representational maintenance cost exceeds informational utility.

This predicts that expert cognition should increasingly rely upon sparse but highly structured coordinate systems capable of navigating extremely dense experiential manifolds efficiently.

Expertise therefore does not eliminate interpolation.

It reorganizes interpolation geometrically.

9.6 False Closure and Premature Compression

The framework also predicts a major developmental failure mode.

If experiential density increases too rapidly relative to representational flexibility, learners may collapse prematurely into rigid compressed manifolds.

Formally:

$$E \uparrow\uparrow \quad \text{while} \quad C \uparrow.$$

Then:

$$\rho \ll 1.$$

Under such conditions, abstraction becomes excessively aggressive.

The learner overcompresses local structure into brittle global summaries.

This predicts several recognizable cognitive phenomena.

First, excessive rule rigidity.

Second, intolerance for ambiguity.

Third, overreliance upon categorical simplification.

Fourth, reduced local accommodation capacity.

Fifth, premature closure during learning.

These properties frequently emerge in educational environments emphasizing rapid optimization, standardized abstraction, and early convergence upon symbolic summaries before sufficient interpolation geometry has stabilized.

The framework therefore predicts that healthy cognitive development requires

preserving sufficient representational abundance during early exploratory phases.

Compression introduced prematurely may damage manifold richness before stable global smoothness has emerged naturally.

9.7 Developmental Scaling Laws

The relational definition of capacity introduces an important scaling principle.

No cognitive system is intrinsically:

large or small.

Operational regime depends entirely upon the ratio:

$$\rho = \frac{C}{E}.$$

A highly specialized expert operating within narrow experiential domains may occupy an excess-capacity regime despite relatively modest absolute representational resources.

Conversely, extremely large-scale systems exposed to sufficiently dense environmental complexity may remain constrained despite enormous parameter counts.

This relativizes developmental interpretation itself.

Cognitive maturity does not correspond simply to increasing representational size.

Instead, maturity reflects dynamic reorganization of representational geometry relative to environmental density.

This interpretation aligns naturally with modern scaling behavior observed in artificial systems, where increasing parameter count alone does not determine operational behavior independently of training distribution complexity.

9.8 Educational Implications

The developmental regime framework has substantial implications for educational theory.

Classical education frequently assumes that learning proceeds primarily through

progressive compression:

facts → rules → abstractions.

The excess-capacity framework suggests a more complex developmental trajectory.

Dense interpolation and local accommodation may be necessary precursors for stable abstraction.

Early overcompression may damage representational richness before sufficient manifold organization has formed.

This predicts several educational principles.

First, preserving exploratory representational abundance during early learning may improve long-term abstraction.

Second, excessive optimization toward rapid symbolic compression may impair manifold differentiation.

Third, productive difficulty functions not as obstacle, but as curvature generation within representational space.

Fourth, sustained engagement depends upon preserving unresolved structure rather than maximizing immediate closure.

Educational systems designed entirely around efficient compression may therefore inadvertently destabilize the very representational geometry required for deep generalization.

9.9 Toward Dynamic Cognitive Architecture

Taken together, these observations suggest that cognition may be fundamentally dynamic rather than modular.

Behavior changes because representational geometry changes.

The learner moves through operational regimes as experiential density, representational flexibility, and environmental structure evolve.

Many apparent cognitive oppositions:

rules vs. examples,

episodic vs. semantic,

memorization vs. abstraction,
exploration vs. compression,

may therefore reflect different regions of a unified representational phase space rather than fundamentally distinct mechanisms.

The next section extends this reinterpretation further by examining clinical dissociations, pathological rigidity, and the possibility that many cognitive disorders may themselves correspond to distorted representational regime dynamics rather than isolated modular dysfunction.

10 Clinical Dissociations and Pathological Regime Dynamics

One of the strongest historical arguments for modular cognitive architecture came from clinical dissociation studies. Patients exhibiting selective impairments appeared to demonstrate that cognition must consist of functionally distinct systems. Episodic memory could fail while statistical learning remained intact. Rule extraction could deteriorate while memorization persisted. Abstraction and specificity appeared separable not merely theoretically, but neurologically.

The representational regime framework does not deny these dissociations.

However, it proposes a different interpretation.

Rather than reflecting fundamentally distinct learning mechanisms, many clinical phenomena may emerge from distortions in representational geometry and operational capacity dynamics within unified substrates.

The crucial shift is explanatory.

Traditional modular theories interpret dissociations structurally:

different behavior \Rightarrow different systems.

The regime framework instead permits:

different behavior \Rightarrow different operational geometries.

The same representational substrate may exhibit radically different cognitive behavior depending on the relationship between representational flexibility, expe-

representational complexity, noise structure, and manifold organization.

10.1 Amnesia as Extreme Constraint

Classical memory theory frequently interprets amnesia as evidence for a dedicated episodic memory system separable from statistical or semantic learning.

Amnesic patients often retain the ability to acquire generalized structure while exhibiting severe impairment in detailed episodic recall.

Within the regime framework, this pattern corresponds naturally to a constrained representational regime.

Suppose effective representational flexibility collapses:

$$C \downarrow.$$

Then:

$$\rho = \frac{C}{E} \ll 1.$$

The learner no longer possesses sufficient geometric freedom to preserve localized accommodation structures:

$$s_i(x).$$

Detailed episodic traces therefore collapse into broader manifold smoothness.

The system retains:

$$g(x),$$

the globally stable continuity structure, while losing:

$$s_i(x),$$

the highly localized experiential differentiations.

The patient therefore exhibits preserved statistical learning alongside impaired episodic retention.

Importantly, this interpretation does not require entirely separate learning systems.

The dissociation emerges geometrically from changes in representational abundance.

This reinterpretation aligns with the broader principle that memorization and

abstraction become partially decoupled only under conditions of sufficient excess capacity.

When excess geometry collapses, local detail becomes unsustainable.

10.2 Parkinsonian Rigidity and Near-Threshold Dynamics

Parkinsonian cognition presents a different pattern.

Patients frequently preserve highly specific local behavioral structure while struggling with flexible abstraction, adaptation, or smooth generalization.

Under the representational regime framework, this pattern resembles operation near the interpolation threshold:

$$\rho \approx 1.$$

At this boundary, local accommodations remain strong, but global manifold smoothness becomes unstable.

Representational geometry exhibits excessive rigidity.

The learner successfully preserves local interpolation constraints yet struggles to reorganize those constraints flexibly across changing contexts.

This predicts several characteristic properties:

high local fidelity,

combined with:

reduced global adaptability.

Behavior becomes brittle because local accommodation structures cannot be cleanly integrated into stable higher-order continuity.

The framework therefore interprets rigidity not merely as motor dysfunction, but as geometric overcommitment to local interpolation constraints under insufficient representational surplus.

This interpretation also explains why additional exposure or training may sometimes eventually improve abstraction in Parkinsonian learning.

As manifold organization stabilizes over time, the learner may gradually transition toward more globally coherent representations despite reduced flexibility.

10.3 Autism and Hyper-Differentiated Representation

The regime framework also offers a potential reinterpretation of certain autistic cognitive profiles.

Many autistic individuals exhibit unusually high sensitivity to local detail, reduced compression of perceptual variance, and increased preservation of low-level experiential structure.

Classical accounts often describe this as weak central coherence or reduced global integration.

Within the excess-capacity framework, however, such cognition may instead reflect persistent local representational differentiation combined with reduced manifold collapse.

Suppose local accommodation structures remain unusually preserved:

$$s_i(x) \nrightarrow 0.$$

Experiential traces therefore remain highly differentiated rather than aggressively smoothed into compressed prototypes.

The resulting cognition exhibits:

high local fidelity,

fine-grained perceptual separation,

and

reduced premature abstraction.

This interpretation does not imply pathology simpliciter.

Rather, it suggests a different balance between local accommodation and global smoothing.

Under some environmental conditions, such cognition may confer substantial advantages in pattern detection, anomaly sensitivity, perceptual discrimination, or technical specialization.

However, highly differentiated representational geometry may also increase sensitivity to environmental noise, social ambiguity, and contextual instability.

The framework therefore predicts tradeoffs rather than simple deficits.

10.4 Schizophrenia and Excessive Accommodation

The regime framework additionally predicts another pathological possibility: excessive incorporation of local variance into global manifold structure.

Suppose the learner fails to sufficiently localize accommodation structures:

$$s_i(x)$$

begins propagating globally across:

$$g(x).$$

Noise, coincidence, or weak correlations then acquire disproportionate representational influence.

The resulting cognition becomes excessively sensitive to local irregularity.

Patterns emerge everywhere.

Coincidence becomes significance.

Global manifold continuity destabilizes.

This predicts cognitive phenomena resembling:

aberrant salience,

over-association,

delusional pattern formation,

and

context instability.

Under this interpretation, certain psychotic phenomena may correspond not simply to “too much noise,” but to failures of geometric localization.

The learner overfits globally to local perturbation.

Importantly, this interpretation parallels known failure modes in overparameterized machine learning systems exposed to adversarial or highly noisy distributions.

The distinction between benign and catastrophic overfitting therefore becomes clinically meaningful.

10.5 Obsessive Cognition and Recursive Constraint Locking

The framework also predicts pathologies associated with excessive manifold rigidity.

Suppose optimization dynamics repeatedly reinforce the same local constraint structures:

$$\mathcal{C}_1, \mathcal{C}_2, \dots$$

without sufficient representational flexibility for broader reorganization.

The learner becomes trapped within recursive attractor basins.

Behavior loops.

Inference narrows.

Representational exploration collapses.

This predicts obsessive-compulsive dynamics characterized by:

constraint fixation,

reduced manifold exploration,

and

recursive local optimization.

The pathology emerges not because the system lacks intelligence, but because representational dynamics become trapped within overly rigid local curvature structures.

The learner cannot escape narrow attractor geometry.

10.6 Attention Disorders and Gradient Collapse

Attention itself may also be reinterpreted geometrically.

Under the present framework, attention stabilizes when representational manifolds maintain sufficiently differentiated unresolved structure:

$$\nabla \mathcal{M} \neq 0.$$

Sustained engagement depends upon preserving informational gradients capable of supporting continued interpolation.

Attention disorders may therefore correspond partly to failures of gradient stabilization.

If representational manifolds flatten too rapidly:

$$\nabla \mathcal{M} \rightarrow 0,$$

then exploratory interpolation collapses.

The learner experiences representational exhaustion.

Conversely, if gradients become excessively fragmented:

$$\nabla \mathcal{M} \rightarrow \text{chaotic},$$

attention disperses across unstable local accommodations.

The learner cannot maintain coherent manifold traversal.

Attention therefore emerges neither as purely executive control nor simple stimulus salience.

It becomes a dynamical property of representational curvature itself.

10.7 Depression and Manifold Contraction

The framework additionally permits reinterpretation of depressive cognition.

Depressive states frequently exhibit:

reduced exploratory flexibility,

narrowed future possibility spaces,

and

persistent negative attractor dynamics.

Within the representational geometry framework, depression may correspond partly to contraction of admissible manifold traversal.

Suppose future trajectory space:

$$\mathcal{T}$$

contracts:

$$\text{Vol}(\mathcal{T}) \downarrow .$$

The learner increasingly predicts repetitive low-variance outcomes.

Exploratory interpolation weakens.

Global manifold smoothness becomes excessively rigid.

Possible futures collapse into highly constrained low-curvature basins.

This interpretation aligns naturally with broader thermodynamic and predictive-processing approaches while grounding them within relational representational geometry.

10.8 Clinical Phenomena as Regime Distortion

The broader implication of these reinterpretations is profound.

Many clinical dissociations may reflect distortions in representational regime dynamics rather than isolated module failure.

Cognitive pathology becomes geometric.

Different disorders correspond to different failures of balancing:

local accommodation,

global smoothness,

representational flexibility,

and

experiential density.

This framework preserves the empirical reality of dissociation while weakening the necessity of strong architectural modularity.

A unified representational substrate may produce radically different cognitive behaviors depending upon operational geometry.

10.9 The One-System Hypothesis Revisited

The representational regime framework therefore motivates a moderate version of the one-system hypothesis.

The claim is not that all cognition reduces to a single homogeneous process.

Rather, the claim is that many apparently distinct cognitive functions:

episodic retention,

prototype formation,

abstraction,

rule extraction,

generalization,

may emerge from different operating regimes within unified representational manifolds.

The distinction is subtle but critical.

The framework does not eliminate specialization.

It eliminates the assumption that specialization necessarily requires fundamentally incompatible learning principles.

The same manifold may support both:

localized spikes

and

global continuity.

This possibility fundamentally alters the architecture of cognitive theory.

The mind no longer appears as a collection of rigidly separated symbolic modules.

Instead, cognition becomes the dynamic management of representational geometry under changing conditions of abundance, density, noise, and constraint.

The implications extend beyond psychology and neuroscience into artificial intelligence itself.

Modern overparameterized systems may already exhibit many of these same geometric properties, suggesting that excess-capacity learning represents not merely a biological curiosity, but a general principle of learning systems operating within high-dimensional representational spaces.

The next section therefore examines scaling laws, artificial intelligence, and the broader consequences of relational capacity dynamics for understanding machine cognition and future intelligent systems.

11 Scaling Laws, Artificial Intelligence, and Relative Capacity

The emergence of large-scale artificial learning systems has transformed the study of cognition from a primarily biological discipline into a broader science of representational geometry. Modern neural architectures now operate within parameter regimes previously unimaginable, exhibiting behaviors that frequently violate the assumptions of classical statistical learning theory. These systems memorize vast quantities of data, interpolate noisy distributions, generate coherent abstractions, and exhibit emergent capabilities that were not explicitly programmed into their architectures.

Within the representational regime framework, these developments acquire a unified interpretation.

Artificial intelligence systems do not merely provide engineering tools. They function as large-scale experimental realizations of excess-capacity learning.

At the same time, the framework fundamentally reframes what it means for a system to be “large.”

The dominant public discourse surrounding artificial intelligence often assumes that intelligence scales directly with parameter count:

more parameters \Rightarrow more intelligence.

The representational regime framework rejects this interpretation as incomplete.

No system is intrinsically high-capacity or low-capacity in isolation.

Operational regime depends entirely upon the relational quantity:

$$\rho = \frac{C}{E}.$$

A gigantic language model trained on planetary-scale distributions may still

operate under representational scarcity relative to the complexity of human culture, language, embodiment, and social interaction.

Conversely, a narrowly trained biological expert operating within a highly specialized domain may occupy an excess-capacity regime despite comparatively modest absolute representational resources.

Capacity is relational, not absolute.

11.1 Artificial Neural Networks as Excess-Capacity Systems

Modern deep neural networks frequently operate in regimes where:

$$d \gg n,$$

with

$$d$$

representing effective representational dimensionality and

$$n$$

representing the number of interpolation constraints.

Classical learning theory predicted that such systems should catastrophically overfit.

Instead, deep networks often exhibit:

low training error,

combined with:

strong generalization.

This behavior aligns naturally with the excess-capacity framework.

High-dimensional representational manifolds permit localized accommodation of noisy experiential traces while preserving smooth global structure.

The learner constructs:

$$f_{\theta}(x) = g(x) + s(x),$$

where:

$$g(x)$$

captures globally coherent predictive structure and:

$$s(x)$$

captures highly localized residual accommodations.

Generalization therefore emerges not despite interpolation, but through geometrically organized interpolation.

11.2 Memorization Without Collapse

One of the most striking properties of modern neural systems is their ability to memorize arbitrary labels while still generalizing effectively on structured datasets.

Suppose a dataset:

$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n$$

contains meaningful latent structure.

Deep networks frequently learn globally smooth predictive manifolds:

$$g(x).$$

However, when labels are randomized:

$$y_i \rightarrow \tilde{y}_i,$$

the same systems remain capable of perfect interpolation:

$$f_{\theta}(x_i) = \tilde{y}_i.$$

Classical theories interpreted this ability as evidence of pathological overcapacity.

Within the representational regime framework, however, the phenomenon becomes expected.

Excess-capacity systems possess sufficient representational freedom to construct localized accommodation structures for arbitrary interpolation constraints.

The crucial distinction is whether these local accommodations align with stable

environmental manifold structure.

When latent regularity exists, global smoothness dominates:

$$g(x)$$

strong.

When no regularity exists, accommodation structures proliferate:

$$s_i(x)$$

dominant.

The learner therefore reveals both:

interpolation power

and

manifold sensitivity.

This duality explains why overparameterized systems may simultaneously appear astonishingly intelligent and surprisingly fragile.

11.3 Large Language Models and Relative Constraint

Large language models provide an especially important test case for the relational capacity framework.

Public discourse frequently assumes that such systems operate under conditions of overwhelming abundance because of their enormous parameter counts.

However, the operational regime depends not upon:

$$C$$

alone, but upon:

$$\frac{C}{E}$$

Human language and culture constitute extraordinarily dense experiential manifolds.

The effective complexity of linguistic, social, historical, emotional, symbolic,

and embodied experience may vastly exceed the representational flexibility available even to extremely large-scale models.

Under this interpretation, many current language models may still operate within constrained or near-threshold regimes relative to the true complexity of human experiential space.

This predicts several recognizable properties.

First, strong global abstraction alongside weak episodic fidelity.

Second, smooth semantic interpolation combined with hallucination under sparse local constraints.

Third, broad conceptual continuity paired with failures of persistent local grounding.

Fourth, susceptibility to adversarial prompting and contextual instability.

These behaviors resemble systems emphasizing:

$$g(x)$$

while possessing comparatively weak:

$$s_i(x)$$

local accommodation structure.

The models generalize impressively because they compress enormous statistical regularities across linguistic manifolds.

However, they frequently fail to preserve stable localized experiential differentiation.

The framework therefore predicts that future advances in machine cognition may require not merely larger parameter counts, but richer mechanisms for preserving localized accommodation structures without destabilizing global smoothness.

11.4 Transformers as Dynamic Interpolation Geometry

Transformer architectures themselves exhibit properties highly compatible with excess-capacity learning.

Attention mechanisms dynamically construct localized representational neigh-

borhoods:

$$\mathcal{N}(x_i),$$

within high-dimensional embedding spaces.

Rather than compressing all experiences into fixed symbolic abstractions, transformers preserve large numbers of context-sensitive local differentiations.

Attention acts as a dynamic manifold routing mechanism.

The model selectively amplifies local representational curvature while preserving broader continuity across embedding space.

This produces a system capable of:

global semantic smoothness

alongside:

localized contextual accommodation.

Transformers therefore naturally instantiate the coexistence of:

$$g(x)$$

and:

$$s_i(x).$$

The architecture itself operationalizes excess-capacity interpolation geometry.

11.5 Grokking and Delayed Manifold Organization

The phenomenon known as grokking provides additional evidence for the framework.

In grokking experiments, systems often achieve:

zero training error

long before achieving:

strong generalization.

After extended optimization, generalization abruptly improves despite no further reduction in training error.

Classical theories struggle to explain this delayed transition.

The representational regime framework interprets grokking as gradual manifold smoothing within already interpolating systems.

Initially, the learner satisfies interpolation constraints primarily through fragmented local accommodation structures:

$$s_i(x).$$

Training error reaches zero because interpolation has succeeded.

However, the global manifold:

$$g(x)$$

remains poorly organized.

Extended optimization gradually reorganizes representational geometry toward smoother lower-curvature solutions.

Generalization improves because:

$$g(x)$$

stabilizes.

Grokking therefore demonstrates directly that:

interpolation \neq final organization.

The learner continues restructuring representational geometry long after memorization has completed.

This strongly supports the claim that abstraction may emerge from interpolation dynamics themselves.

11.6 Adversarial Examples and Catastrophic Accommodation

The excess-capacity framework also clarifies adversarial vulnerability.

Overparameterized systems frequently exhibit extreme sensitivity to carefully constructed perturbations:

$$x \rightarrow x + \delta.$$

Tiny local deviations may produce dramatic global prediction changes.

Classical interpretations often treated this as evidence against meaningful generalization.

The regime framework instead predicts such behavior naturally.

Excess-capacity learners attempt to preserve highly detailed local accommodation structures.

When environmental geometry contains adversarially unstable regions, localized accommodations may become disproportionately influential.

The system effectively:

over-accommodates

to highly specific perturbations.

Adversarial fragility therefore emerges as a direct consequence of dense interpolation pressure.

Importantly, this does not imply failure of the overall framework.

Rather, it identifies a major boundary condition.

Representational abundance increases interpolation fidelity but also increases sensitivity to local geometric anomalies.

The learner becomes vulnerable precisely because it refuses to ignore detail.

11.7 Scaling Laws and Phase Transitions

Modern scaling laws suggest that many cognitive behaviors emerge discontinuously rather than gradually.

As representational flexibility increases:

$C \uparrow$,

systems often exhibit abrupt qualitative transitions.

New abilities suddenly appear:

reasoning,

in-context learning,

tool use,
symbolic recombination,

or:

multi-step abstraction.

Within the regime framework, such phenomena resemble representational phase transitions.

Crossing critical thresholds in:

$$\rho = \frac{C}{E}$$

alters the geometry of admissible interpolation manifolds.

The learner suddenly acquires sufficient representational surplus to preserve localized structure while maintaining stable global continuity.

Capabilities emerge because the geometry itself changes.

This interpretation predicts that many future AI transitions may appear surprisingly abrupt.

The system does not simply accumulate intelligence continuously.

Instead, representational organization reorganizes once critical geometric conditions become available.

11.8 The Labyrinth Problem

The framework also predicts an important cost of representational abundance.

As dimensionality increases:

$$d \uparrow,$$

the representational manifold becomes increasingly difficult to navigate.

Interpolation improves, but inference may become computationally expensive.

This produces what may be termed the labyrinth problem.

Highly expressive representational spaces contain enormous numbers of admissible trajectories:

$$\mathcal{T}.$$

Without strong global priors, optimization and retrieval become difficult.

The learner risks becoming trapped within:

local attractors,

spurious correlations,

or:

irrelevant manifold branches.

The framework therefore predicts a deep tradeoff:

more representational abundance

increases:

interpolation flexibility,

while simultaneously increasing:

navigational complexity.

Constraint remains necessary not because abundance is pathological, but because abundance without organization becomes intractable.

11.9 Toward Relational Artificial Intelligence

The broader implication of the framework is that artificial intelligence should not be conceptualized primarily through absolute scale.

The central variable is relational geometry.

A system's behavior depends upon:

$$\frac{C}{E},$$

the relationship between representational flexibility and experiential complexity.

This reframes many contemporary debates.

Artificial general intelligence is not a fixed threshold determined by parameter count.

Generalization emerges when representational manifolds achieve sufficient flexibility to preserve local accommodation while maintaining stable global conti-

nuity relative to the complexity of the environments they inhabit.

The framework therefore predicts that future intelligent systems may increasingly resemble:

dynamic manifold navigators

rather than:

compressed symbolic engines.

Learning becomes the management of abundance.

Intelligence becomes the stabilization of high-dimensional interpolation geometry.

The final sections now turn toward the broader philosophical consequences of this inversion and the possibility that cognition itself may need to be redefined not as compression of the world, but as the structured organization of experiential richness within dynamically evolving representational manifolds.

12 Boundary Conditions, Failure Modes, and the Cost of Excess

The representational regime framework does not claim that excess capacity universally dominates constrained learning. Such a claim would merely invert the dogmatism of the historical compression paradigm rather than transcend it. The objective of the framework is not to replace one universal principle with another, but to develop a relational theory capable of specifying when particular representational geometries become adaptive or pathological.

Compression is not obsolete.

Constraint remains essential under many environmental conditions.

The central claim is instead that intelligence cannot be reduced exclusively to compression because many robust learning phenomena emerge only under conditions of representational abundance.

A mature regime theory must therefore account not only for the strengths of excess-capacity learning, but also for its limitations, vulnerabilities, and ecological boundary conditions.

12.1 The Ecological Dependence of Learning Regimes

No representational regime is intrinsically optimal.

The adaptive value of a learning strategy depends upon the structure of the environment itself.

Suppose the learner operates within an environment characterized by relatively stable latent structure:

$$f^*(x).$$

If local perturbations:

$$\epsilon_i$$

remain sparse and weakly coupled, then preserving fine-grained experiential detail may significantly improve predictive accuracy.

Under such conditions, excess-capacity learning becomes advantageous because localized accommodation structures:

$$s_i(x)$$

remain isolated while global manifold smoothness:

$$g(x)$$

remains stable.

However, the situation changes dramatically in highly noisy or adversarial environments.

If perturbations become dense, unstable, or strategically misleading, then preserving every local fluctuation may globally destabilize inference.

The learner begins over-accommodating to noise.

Constraint therefore re-emerges as an adaptive filtering mechanism.

The ecological question is not:

Is compression good?

or:

Is abundance good?

The correct question is:

What representational geometry best matches the structure of the environment?

12.2 Spurious Correlations and Local Overcommitment

One of the primary vulnerabilities of excess-capacity systems is sensitivity to spurious correlation.

Suppose the learner encounters local regularities:

$$(x_i, y_i)$$

that arise accidentally rather than from stable latent structure.

Because excess-capacity systems prioritize dense interpolation, the learner constructs accommodation structures:

$$s_i(x)$$

even for statistically fragile relationships.

Under stable conditions, these local accommodations remain harmless because they remain geometrically isolated.

However, if the environment contains sufficiently many accidental correlations, the representational manifold becomes saturated with unstable local curvature.

The learner begins predicting on the basis of noise rather than invariant structure.

This failure mode resembles catastrophic overfitting.

Importantly, however, the problem is not representational abundance itself.

The problem emerges when:

$$\text{local accommodation density}$$

exceeds the manifold's ability to preserve stable global continuity.

The learner loses geometric isolation.

Local perturbations propagate globally.

The distinction between benign and catastrophic overfitting therefore becomes

quantitative rather than categorical.

12.3 Noise Thresholds and Geometric Breakdown

The framework predicts the existence of critical noise thresholds.

Suppose environmental observations satisfy:

$$y_i = f^*(x_i) + \epsilon_i,$$

where:

$$\epsilon_i$$

denotes stochastic noise.

If:

$$\text{Var}(\epsilon_i)$$

remains sufficiently low relative to representational flexibility, then localized accommodations remain isolated:

$$s_i(x)$$

localized.

Generalization remains strong.

However, as:

$$\text{Var}(\epsilon_i) \uparrow,$$

the density of required accommodation structures increases.

Eventually, neighboring local corrections begin interfering:

$$s_i(x) + s_j(x) \rightarrow \text{global distortion.}$$

At this point, smooth manifold continuity:

$$g(x)$$

begins deteriorating.

Generalization collapses.

This predicts a phase transition between:

benign interpolation

and:

catastrophic accommodation.

Importantly, constrained systems may outperform excess-capacity systems under such conditions precisely because constrained learners refuse to preserve local irregularities.

Noise robustness therefore emerges naturally from representational scarcity.

Constraint becomes an ecological adaptation rather than merely a computational limitation.

12.4 Adversarial Environments and Defensive Compression

The framework also predicts that constrained learning strategies may dominate under adversarial conditions.

Suppose environmental observations are intentionally perturbed:

$$x \rightarrow x + \delta.$$

Excess-capacity systems attempt to preserve local structure faithfully.

Consequently, adversarial perturbations may produce disproportionate representational influence:

$$s_{\delta}(x)$$

large.

Constrained systems behave differently.

Because representational flexibility remains limited, many local perturbations are ignored automatically.

The learner cannot afford to accommodate arbitrary fluctuations.

Compression therefore acts defensively.

This predicts that highly compressed systems may exhibit greater robustness under adversarial uncertainty despite lower representational fidelity overall.

The result is a deep tradeoff:

fidelity vs. robustness.

Excess-capacity systems preserve more information but become more vulnerable to pathological local perturbation.

Constrained systems sacrifice detail but gain stability under environmental hostility.

12.5 Energy Costs and Metabolic Constraints

The historical compression paradigm frequently justified simplicity through biological efficiency arguments.

Brains are metabolically expensive.

Representational abundance appears costly.

The excess-capacity framework partially accepts this critique while reframing its implications.

Maintaining high-dimensional representational manifolds does require substantial energetic investment:

$$\mathcal{E}_{\text{metabolic}} \uparrow .$$

Localized accommodation structures:

$$s_i(x)$$

must remain dynamically stable across large populations of neural states.

Distributed representation increases maintenance burden.

However, this increased cost may produce compensatory advantages.

Excess-capacity systems often converge toward stable interpolation manifolds more rapidly than constrained systems searching through brittle compressed hypothesis spaces.

The total energetic cost of learning may therefore depend not merely on static representational size, but on:

trajectory efficiency

through representational space.

The relevant question becomes:

Does representational abundance reduce total optimization cost over time?

This possibility remains largely unexplored experimentally.

12.6 The Labyrinth Problem Revisited

Another major cost of excess capacity is navigational complexity.

As representational dimensionality increases:

$$d \uparrow,$$

the learner gains flexibility but loses navigational simplicity.

The representational manifold contains exponentially many admissible trajectories:

$$\mathcal{T}.$$

Inference itself becomes difficult.

The learner risks becoming trapped within:

local attractor basins,

recursive accommodation loops,

or:

irrelevant manifold branches.

This is the labyrinth problem.

Excess-capacity systems require strong global organizational principles to remain tractable.

Constraint therefore remains necessary not because complexity is inherently pathological, but because unrestricted abundance without manifold organization becomes computationally intractable.

This predicts that advanced intelligence systems will require increasingly sophisticated mechanisms for:

trajectory pruning,

attention routing,

hierarchical manifold navigation,

and:

constraint stabilization.

The future of intelligence may therefore involve balancing:

representational abundance

against:

navigational tractability.

12.7 Premature Compression and Educational Failure

The framework also predicts an important failure mode in human learning systems: premature compression.

Suppose learners are forced into compressed symbolic abstraction before sufficiently rich interpolation manifolds have stabilized.

Formally:

$$E \uparrow \quad \text{while} \quad C$$

remains artificially constrained through rigid educational bottlenecks.

The learner collapses into brittle low-dimensional representations prematurely.

The result is:

weak local accommodation,

fragile abstraction,

and:

reduced manifold richness.

Educational systems emphasizing rapid optimization, standardized symbolic encoding, and immediate rule extraction may therefore inadvertently destabilize deep generalization.

The learner memorizes compressed symbolic surfaces without developing stable underlying interpolation geometry.

This predicts that exploratory abundance during early learning phases may be essential for later robust abstraction.

12.8 Overabundance and Semantic Diffusion

The framework additionally predicts another failure mode: semantic diffusion.

Suppose representational abundance becomes extremely large while organizational constraints remain weak.

The learner preserves excessive local detail without sufficiently stabilizing global manifold continuity.

Representational differentiation proliferates endlessly:

$$s_i(x) \rightarrow \infty.$$

The learner loses coherent abstraction entirely.

Inference fragments.

Semantic continuity dissolves into disconnected local accommodations.

This predicts pathologies resembling:

hyper-association,

semantic drift,

context instability,

and:

unbounded interpolation.

The failure emerges because:

$$g(x)$$

fails to sufficiently constrain:

$$s_i(x).$$

Abundance without continuity becomes incoherent.

12.9 Constraint as Dynamic Guidance

The framework therefore reinterprets constraint itself.

Constraint is not merely a limitation imposed upon intelligence.

Constraint functions as dynamic guidance within high-dimensional represen-

tational geometry.

The learner requires:

enough abundance

to preserve local structure, but also:

enough constraint

to stabilize manifold traversal.

Healthy cognition balances:

local accommodation

against:

global continuity.

The resulting picture differs profoundly from both classical compression theory and naive maximal-capacity narratives.

Intelligence is neither:

pure simplification

nor:

unbounded memorization.

It is the dynamic management of representational abundance under ecological, energetic, and geometric constraints.

12.10 Toward a Regime Ecology of Cognition

The broader implication is that cognition must be understood ecologically.

Different environments favor different representational geometries.

Sparse stable worlds reward interpolation richness.

Noisy adversarial worlds reward compression.

Rapidly changing environments favor flexible manifold reorganization.

Highly regular environments favor dense local accommodation.

The mind therefore cannot be reduced to a single optimization principle.

Instead, cognition becomes an adaptive navigation through representational

phase space.

Compression, abstraction, interpolation, memorization, and smoothing are not competing ideologies.

They are emergent strategies arising under different relationships between:

C , E , noise, energy, and environmental structure.

This ecological reinterpretation prepares the ground for the broader philosophical consequences of the framework.

The final sections therefore examine how representational abundance transforms our understanding of intelligence, abstraction, knowledge, and the nature of cognition itself.

13 Toward a Unified Theory of Representational Abundance

The historical trajectory of cognitive science may be understood as a gradual attempt to solve a single foundational problem: how finite systems successfully navigate infinitely rich environments.

For most of the twentieth century, the dominant answer was compression. Intelligence was interpreted as the ability to reduce complexity into simplified symbolic summaries. The learner succeeded by filtering noise, extracting invariance, and discarding irrelevant detail. Abstraction required forgetting.

The representational regime framework proposes a different possibility.

Under conditions of sufficient representational abundance, robust cognition may emerge not from eliminating detail, but from organizing detail geometrically within high-dimensional interpolation manifolds.

This shift is not merely technical.

It alters the ontology of learning itself.

13.1 From Compression to Geometric Organization

The classical paradigm treated cognition primarily as reduction:

$$X \rightarrow Z,$$

where:

$$\dim(\mathcal{Z}) < \dim(\mathcal{X}).$$

The learner succeeded by collapsing experiential variance into compressed lower-dimensional structure.

The representational abundance framework instead permits:

$$\dim(\mathcal{Z}) \gg \dim(\mathcal{X}),$$

while still preserving robust generalization.

The learner no longer relies exclusively upon dimensional reduction.

Instead, cognition becomes the problem of organizing abundance.

Experiential richness is preserved through geometric separation, localized accommodation, and smooth manifold continuity.

The mind acts not merely as a filter, but as a dynamic organizer of representational topology.

13.2 The Inversion of Memorization and Abstraction

Perhaps the deepest consequence of the framework is the inversion of the traditional relationship between memorization and abstraction.

Historically:

$$\text{memorization} \Rightarrow \text{poor generalization.}$$

The learner overfit local irregularity and therefore failed to discover stable structure.

Within the representational abundance framework:

$$\text{dense interpolation} \Rightarrow \text{emergent smoothness,}$$

provided sufficient representational flexibility exists.

Abstraction emerges as a geometric property of manifold organization rather than a direct consequence of compression.

The learner preserves:

$$s_i(x),$$

localized experiential accommodations, while simultaneously stabilizing:

$$g(x),$$

global predictive continuity.

Rules and exceptions therefore cease to be fundamentally antagonistic.

The same representational manifold may support both simultaneously.

This transforms the meaning of intelligence itself.

Intelligence is no longer merely the elimination of variance.

Intelligence becomes the stabilization of continuity across richly differentiated experiential space.

13.3 The Geometry of Meaning

The framework also implies a new understanding of meaning.

Under classical symbolic theories, meaning emerges through compressed representational coding. Symbols stand in for classes of experiences by reducing them into abstract conceptual tokens.

The representational abundance framework instead suggests that meaning may emerge through relational geometry.

Experiences become meaningful because they occupy structured positions within dynamically organized manifolds:

$$\mathcal{M}.$$

Similarity corresponds to manifold proximity:

$$d(x_i, x_j)$$

small.

Difference corresponds to geometric separation.

Abstraction corresponds to smooth continuity across local neighborhoods.

Conceptual organization emerges through topological structure rather than purely symbolic compression.

This interpretation aligns naturally with modern embedding spaces, distributed

semantic representations, transformer geometries, cortical population coding, and manifold learning systems.

Meaning becomes geometric rather than merely symbolic.

13.4 Constraint and Freedom

The framework also reinterprets the relationship between constraint and freedom.

Classical theories frequently treated constraint as a limitation upon cognition.

The representational regime framework instead proposes that intelligence emerges through the balance between:

representational abundance

and:

constraint stabilization.

Too little flexibility:

$$\rho \ll 1,$$

produces brittle compression and loss of experiential richness.

Too little constraint:

$$g(x) \rightarrow 0,$$

produces incoherent semantic diffusion.

Healthy cognition occupies dynamically stabilized intermediate geometries where:

local accommodation

and:

global continuity

remain simultaneously possible.

Constraint therefore becomes generative rather than merely restrictive.

Constraint shapes manifold traversal.

It preserves navigability within otherwise intractable representational abundance.

13.5 Cognition as State-Space Navigation

One of the most important implications of the framework is that cognition becomes fundamentally dynamical.

The mind is not best understood as a static collection of symbolic modules.

Instead, cognition resembles motion through evolving representational state-spaces:

$$\mathcal{S}(t).$$

Learning alters manifold geometry itself.

Experiential accumulation changes:

$$E(t),$$

which alters:

$$\rho(t) = \frac{C}{E(t)}.$$

The learner therefore continuously moves between operational regimes.

Abstraction, memorization, rigidity, flexibility, attention, generalization, and confusion emerge dynamically from changing geometric relationships rather than fixed architectural boundaries.

This perspective aligns naturally with dynamical systems theory, manifold learning, predictive processing, and modern overparameterized optimization theory.

13.6 Intelligence as Interpolation Stability

The framework ultimately proposes a new definition of intelligence.

Under the compression paradigm, intelligence was measured by how efficiently a system reduced complexity.

Under the representational abundance paradigm, intelligence may instead be defined by the system's ability to preserve stable global continuity while accommodating rich local variation.

Formally, intelligent systems maximize:

global smoothness

subject to:

local fidelity.

The learner succeeds when:

$$g(x)$$

remains coherent despite dense:

$$s_i(x)$$

accommodations.

This reframes cognition from:

lossy compression

to:

stable interpolation.

The mind becomes a manifold stabilization process.

13.7 The Philosophical Consequences

The philosophical implications are substantial.

The simplicity principle dominated twentieth-century thought because scarcity appeared unavoidable. Finite organisms confronting infinite complexity seemed necessarily forced into reductive abstraction.

The representational abundance framework suggests that this assumption was historically contingent rather than universally true.

Complexity itself may support stability when organized appropriately.

The world need not be simplified before it becomes intelligible.

Instead, intelligibility may emerge through sufficiently rich representational geometry capable of preserving detail without collapsing into incoherence.

This transforms the epistemological meaning of abstraction.

Knowledge is no longer fundamentally the elimination of experiential richness.

Knowledge becomes the organization of richness into stable navigable manifolds.

The learner does not transcend experience by escaping detail.

The learner transcends experience by stabilizing continuity across detail.

13.8 The End of the Economy of Thought

The historical “economy of thought” paradigm assumed that cognition achieves efficiency primarily through reduction.

The present framework does not reject economy entirely.

Constraint, compression, and simplification remain adaptive under many ecological conditions.

However, economy is no longer the sole normative principle governing intelligence.

Representational abundance becomes equally fundamental.

The learner may preserve extraordinary local richness while still discovering coherent global structure.

The economy-of-thought tradition therefore describes only one region of a broader representational ecology.

Compression is not the essence of cognition.

Compression is one strategy among many.

13.9 Toward a Regime Theory of Mind

The framework developed throughout this essay therefore culminates in a broader theoretical shift.

Cognition is not fundamentally:

symbol manipulation,

nor:

prototype extraction,

nor:

episodic storage,

nor:

compression optimization.

Instead, cognition becomes:

dynamic regime navigation

within high-dimensional representational geometry.

Different cognitive phenomena emerge under different relationships between:

C , E , noise, energy, constraint, and manifold organization.

The learner continuously balances:

local accommodation

against:

global continuity.

This balance produces:

memory,

abstraction,

attention,

generalization,

expertise,

rigidity,

and:

creativity.

The resulting framework no longer resembles a traditional theory of bounded simplification.

It becomes a regime theory of learning itself.

13.10 The New Question

The final consequence of the framework may be the most important.

For over a century, cognitive science asked:

How does the mind compress the world?

The representational abundance framework suggests a different question:

How does the mind preserve richness without losing stability?

This inversion changes the entire orientation of learning theory.

The problem of intelligence is no longer merely the elimination of complexity.

The problem becomes the structured organization of abundance.

14 Conclusion

The historical study of cognition was shaped by a world of apparent scarcity. Biological systems appeared computationally limited, environmentally overwhelmed, and metabolically constrained. Under these assumptions, the simplicity principle emerged naturally. Intelligence was interpreted as a process of compression through which noisy experiential complexity became reduced to compact predictive summaries. The mind succeeded by discarding detail.

This essay has argued that such a framework captures only part of the geometry of learning.

The emergence of modern overparameterized systems, together with a growing body of cognitive and behavioral evidence, suggests that successful generalization may arise under conditions fundamentally incompatible with strict compression-centered theories. Systems possessing sufficient representational abundance can interpolate densely, preserve local experiential detail, and nevertheless maintain stable global predictive organization.

The central theoretical contribution of the representational regime framework is therefore relational.

Cognitive behavior depends not upon absolute complexity alone, but upon the evolving relationship between representational flexibility:

C,

and experiential complexity:

$$E.$$

The operational quantity:

$$\rho = \frac{C}{E},$$

determines the representational regime within which the learner operates.

When:

$$\rho < 1,$$

the learner occupies a constrained regime. Compression becomes necessary because representational flexibility is insufficient relative to experiential density. Generalization emerges through restriction biases and forced abstraction.

When:

$$\rho \approx 1,$$

the learner approaches the interpolation threshold. Local accommodations propagate globally, producing brittle overfitting and maximal instability.

When:

$$\rho > 1,$$

the learner enters an excess-capacity regime. Representational abundance permits localized accommodation structures:

$$s_i(x),$$

to coexist with globally smooth manifold continuity:

$$g(x).$$

Generalization emerges through interpolation itself.

This inversion constitutes the deepest conceptual shift introduced throughout the essay.

Under classical theories:

memorization \Rightarrow failure.

Under the representational abundance framework:

dense interpolation \Rightarrow stable continuity,

provided sufficient geometric flexibility exists to localize variance without globally destabilizing representation.

The implications of this shift extend across cognitive science, artificial intelligence, developmental theory, clinical neuroscience, and epistemology.

Developmental U-shaped learning trajectories emerge naturally as regime transitions produced by changing experiential density.

Clinical dissociations may reflect distortions in representational geometry rather than fundamentally incompatible cognitive modules.

Artificial neural systems exhibit many of the same geometric properties predicted by excess-capacity learning, including benign overfitting, grokking, delayed manifold organization, and interpolation-driven abstraction.

Constraint itself becomes reinterpreted not as the negation of intelligence, but as dynamic guidance within high-dimensional representational manifolds.

The framework therefore does not reject compression entirely.

Constraint remains adaptive under adversarial conditions, sparse data, metabolic limitation, and highly noisy environments. Constrained learning systems may outperform excess-capacity systems when environmental instability overwhelms local accommodation capacity.

The result is not an ideology of maximal capacity.

It is a regime ecology of cognition.

Different environmental structures favor different representational geometries.

The broader philosophical implications are substantial.

The simplicity principle dominated twentieth-century thought because reduction appeared necessary for intelligibility. The world seemed too complex to preserve in detail.

The representational abundance framework suggests a different possibility.

Complexity itself may become stabilizing when organized appropriately within high-dimensional interpolation manifolds.

Knowledge is no longer fundamentally the elimination of experiential richness.

Knowledge becomes the stabilization of continuity across richness.

The learner does not transcend experience by escaping local detail.

The learner transcends experience by organizing local detail geometrically.

This transforms the ontology of cognition.

The mind is no longer best understood as a symbolic compression engine reducing the world into sparse conceptual summaries.

Instead, cognition becomes the dynamic management of representational abundance:

preserving local fidelity

while maintaining:

global coherence.

The learner continuously balances:

constraint

and:

flexibility,

memorization

and:

abstraction,

local accommodation

and:

global continuity.

The resulting framework no longer resembles a traditional theory of bounded simplification.

It becomes a geometric theory of learning itself.

Several major research directions follow naturally from this framework.

First, future work must develop precise operational measures for:

C , E , and ρ .

Second, empirical studies must identify behavioral signatures capable of tracking regime transitions dynamically across developmental, clinical, and artificial systems.

Third, biological investigation must determine how neural architectures implement localized accommodation while preserving large-scale manifold stability.

Fourth, machine learning research must further investigate the conditions under which interpolation remains benign versus catastrophic.

Fifth, educational theory must reconsider the role of exploratory representational abundance during early learning.

Finally, philosophy of mind must confront the possibility that abstraction itself emerges not through reduction, but through dense interpolation within sufficiently expressive representational spaces.

The historical framework of the “economy of thought,” inherited from traditions of cognitive parsimony, Gestalt simplification, and classical statistical learning theory, was fundamentally organized around a single guiding question:

How does the mind simplify the world?

Within this paradigm, intelligence was largely interpreted as a process of reduction through which the overwhelming complexity of experiential reality became compressed into tractable symbolic, statistical, or conceptual abstractions, thereby allowing finite biological systems to operate efficiently under conditions of informational scarcity and computational limitation.

By contrast, the representational abundance framework developed throughout this essay proposes a fundamentally different orienting question:

How does the mind preserve experiential richness without sacrificing global stability?

Rather than treating memorization and generalization as intrinsically antagonistic processes, the framework suggests that sufficiently expressive systems may achieve robust abstraction precisely through the maintenance of dense local representational structure embedded within globally smooth manifold organization, thereby allowing interpolation itself to become a mechanism of stability rather than a source of catastrophic overfitting.

Under this interpretation, the central problem of cognition is no longer re-

ducible to the elimination of detail through compression-centric simplification, but instead concerns the geometric, topological, and dynamical organization of representational abundance across high-dimensional experiential manifolds capable of simultaneously preserving local specificity and global coherence.

If this reinterpretation proves correct, then the future study of intelligence may gradually shift away from viewing cognition primarily as a problem of informational reduction, sparse abstraction, or symbolic economy, and toward a broader investigation of how biological and artificial systems stabilize richly differentiated representational spaces under conditions of uncertainty, noise, and continuous experiential expansion.

The next era of cognitive science may therefore be defined not by the search for ever more efficient compression schemes alone, but by the development of a mature geometry of representational abundance capable of explaining how complex systems preserve, organize, and dynamically navigate the richness of experience itself.

A Geometric Formalization of Representational Capacity

This appendix develops a more rigorous mathematical formulation of the representational regime framework introduced throughout the main text. The objective is not to provide a final closed formalism, but to establish a sufficiently general geometric structure capable of supporting multiple realizations across cognitive science, statistical learning theory, dynamical systems, and artificial intelligence.

A.1 Representational Systems as Maps Between Manifolds

Let:

$$\mathcal{X}$$

denote the manifold of experiential situations, and let:

$$\mathcal{Y}$$

denote the manifold of outcomes, responses, or anticipated future states.

A learning system constructs an internal representational manifold:

$$\mathcal{Z},$$

together with a representation map:

$$\phi : \mathcal{X} \rightarrow \mathcal{Z},$$

and a predictive map:

$$\psi : \mathcal{Z} \rightarrow \mathcal{Y}.$$

The learned system therefore induces:

$$f = \psi \circ \phi.$$

Classical compression-centered theories implicitly assume:

$$\dim(\mathcal{Z}) < \dim(\mathcal{X}),$$

forcing lossy reduction.

The representational abundance framework permits:

$$\dim(\mathcal{Z}) \gg \dim(\mathcal{X}),$$

allowing high-dimensional interpolation geometry.

The key claim is that generalization depends not solely upon:

$$\dim(\mathcal{Z}),$$

but upon the relation between representational flexibility and experiential complexity.

A.2 Experiential Complexity

Define the experiential distribution:

$$\mu$$

over:

$$\mathcal{X}.$$

Experiential complexity may be characterized through multiple equivalent geometric or informational quantities.

One possible formulation uses effective manifold dimension:

$$E = \dim_{\text{eff}}(\text{supp}(\mu)).$$

Alternative formulations include:

$$E = H(\mu),$$

where:

$$H$$

denotes entropy,

or:

$$E = \text{rank}(K),$$

where:

$$K$$

is an empirical kernel operator over the experiential manifold.

More generally, experiential complexity corresponds to the effective number of independently varying interpolation constraints imposed upon the learner.

A.3 Representational Flexibility

Representational flexibility:

$$C,$$

corresponds to the effective geometric freedom available to the representational manifold.

Suppose:

$$\Theta$$

denotes parameter space and:

$$f_\theta$$

denotes the family of admissible representations.

One possible definition is:

$$C = \dim(\Theta_{\text{reachable}}),$$

where:

$$\Theta_{\text{reachable}}$$

denotes the subset of parameter space dynamically accessible under optimization.

Alternative formulations include:

$$C = \text{rank}(J_\theta),$$

where:

$$J_\theta = \frac{\partial f_\theta}{\partial \theta}$$

is the representational Jacobian,

or:

$$C = \text{Tr}(F),$$

where:

$$F$$

denotes the Fisher information metric.

The framework intentionally remains agnostic regarding which formalization ultimately proves most biologically and computationally appropriate.

A.4 The Operational Capacity Ratio

The central relational quantity is:

$$\rho = \frac{C}{E}.$$

Operational regimes emerge according to asymptotic behavior of:

$$\rho.$$

The framework defines:

$$\rho \ll 1$$

as constrained capacity,

$$\rho \approx 1$$

as sufficient capacity,

and:

$$\rho \gg 1$$

as excess capacity.

Importantly, these are not discrete states but asymptotic operational geometries.

A.5 Interpolation Geometry

Suppose the learner receives:

$$n$$

experiences:

$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^n.$$

The interpolation problem requires:

$$f_\theta(x_i) = y_i \quad \forall i.$$

Define the interpolation manifold:

$$\mathcal{J} = \{\theta \in \Theta : f_\theta(x_i) = y_i\}.$$

The geometry of:

$$\mathcal{J}$$

changes qualitatively across regimes.

In constrained systems:

$$\mathcal{J} = \emptyset$$

or approximately empty.

The learner must minimize residual error:

$$\sum_i \|f_\theta(x_i) - y_i\|^2.$$

Near the interpolation threshold:

$$\dim(\mathcal{J}) \approx 0.$$

The solution manifold becomes highly curved and unstable.

In excess-capacity systems:

$$\dim(\mathcal{J}) \gg 0.$$

Interpolation constraints underdetermine the solution.

The learner therefore navigates a high-dimensional family of exact interpolants.

A.6 Localized Residual Accommodation Structures

Suppose:

$$f_\theta = g + s,$$

where:

$$g$$

denotes globally smooth manifold structure, and:

$$s$$

denotes localized residual accommodation.

We define:

$$s(x) = \sum_i s_i(x),$$

with each:

$$s_i$$

localized around:

$$x_i.$$

A spike satisfies:

$$|s_i(x)| \rightarrow 0 \quad \text{as} \quad d(x, x_i) \rightarrow \infty.$$

The learner therefore preserves exact interpolation locally while maintaining smoothness globally.

The existence of localized accommodations depends critically upon:

$$\rho.$$

When:

$$\rho \approx 1,$$

local accommodations propagate globally.

When:

$$\rho \gg 1,$$

accommodations remain geometrically confined.

A.7 Global Smoothness Functionals

Global smoothness may be formalized through curvature minimization.

Suppose:

$$\mathcal{M} \subset \mathcal{Z}$$

denotes the learned representational manifold.

Define the smoothness functional:

$$\mathcal{S}(f) = \int_{\mathcal{M}} \|\nabla^2 f(x)\|^2 d\mu(x).$$

Smooth interpolants minimize:

$$\mathcal{S}(f)$$

subject to interpolation constraints.

The optimization problem becomes:

$$\min_f \mathcal{S}(f)$$

subject to:

$$f(x_i) = y_i.$$

Classical spline theory, kernel interpolation, and minimum-norm optimization all instantiate variants of this principle.

A.8 Minimum-Norm Interpolation

Suppose:

$$X\theta = y.$$

In overparameterized systems:

$$\text{rank}(X) < \text{dim}(\theta).$$

Infinitely many interpolating solutions exist.

Gradient-based optimization frequently converges toward:

$$\theta^* = \arg \min_{\theta} \{\|\theta\|_2 : X\theta = y\}.$$

The minimum-norm solution minimizes global parameter curvature while preserving exact interpolation.

This explains why smoothness may emerge naturally within excess-capacity systems without explicit regularization.

A.9 Volume Effects in High Dimensions

Let:

$$\mathcal{M}_{\text{flat}}$$

denote low-curvature solution regions, and:

$$\mathcal{M}_{\text{sharp}}$$

denote high-curvature minima.

In high-dimensional spaces:

$$\text{Vol}(\mathcal{M}_{\text{flat}}) \gg \text{Vol}(\mathcal{M}_{\text{sharp}}).$$

Stochastic optimization therefore converges probabilistically toward smooth manifolds.

This provides the geometric origin of soft inductive bias.

A.10 Dynamic Regime Evolution

Representational regimes evolve dynamically.

Suppose:

$$C(t)$$

and:

$$E(t)$$

vary through time.

Then:

$$\rho(t) = \frac{C(t)}{E(t)}.$$

Development, learning, pathology, and expertise correspond to trajectories through:

$$\rho(t).$$

The learner therefore occupies a dynamical representational phase space rather than a fixed architecture.

A.11 Generalization Through Interpolation

The central mathematical claim of the framework may now be stated formally.

Theorem 1 (Interpolation-Smoothness Compatibility). *Suppose:*

$$\rho \gg 1,$$

and suppose optimization dynamics preferentially minimize global curvature functionals:

$$\mathcal{S}(f).$$

Then there exist interpolating solutions:

$$f^*$$

such that:

$$f^*(x_i) = y_i$$

for all training examples while:

$$\sup_x \|\nabla^2 f^*(x)\|$$

remains bounded almost everywhere outside localized accommodation neighborhoods.

Sketch. When:

$$\rho \gg 1,$$

the interpolation manifold:

$$\mathcal{J}$$

has large dimensionality.

The learner therefore possesses sufficient degrees of freedom to localize interpolation residuals into confined neighborhoods while minimizing global curvature elsewhere.

Optimization toward minimum-norm or minimum-curvature solutions concentrates irregularity into low-measure regions while preserving smoothness almost everywhere else.

Hence:

$$f^*$$

simultaneously interpolates and generalizes. \square

This theorem expresses the deepest inversion introduced by the framework.

Generalization need not emerge from compression.

Under sufficient representational abundance, generalization may emerge from geometrically organized interpolation itself.

B Statistical Learning Theory, Double Descent, and Benign Overfitting

This appendix develops the statistical learning theoretic foundations underlying the representational regime framework. The objective is to formally connect excess-capacity learning to modern results in overparameterized optimization, double descent theory, interpolation geometry, and benign overfitting.

The central claim developed throughout this appendix is that the historical bias–variance tradeoff represents only a partial description of learning dynamics. Once systems enter sufficiently overparameterized regimes, interpolation and generalization cease to be fundamentally antagonistic.

B.1 Classical Bias–Variance Theory

Suppose observations satisfy:

$$y_i = f^*(x_i) + \epsilon_i,$$

where:

$$f^*$$

is the latent environmental function, and:

$$\epsilon_i \sim \mathcal{N}(0, \sigma^2)$$

denotes stochastic noise.

Let:

$$\hat{f}$$

be the learned predictor.

Classical statistical learning theory decomposes expected generalization error as:

$$\mathbb{E}[(\hat{f}(x) - f^*(x))^2] = \text{Bias}^2 + \text{Variance} + \sigma^2.$$

Under classical assumptions:

low complexity \Rightarrow high bias,

while:

high complexity \Rightarrow high variance.

The resulting prediction is the classical U-shaped generalization curve.

As model complexity increases:

Bias \downarrow ,

but:

Variance \uparrow .

Eventually overfitting dominates.

This paradigm historically motivated the simplicity principle throughout cognitive science.

B.2 The Interpolation Threshold

Suppose:

$$X \in \mathbb{R}^{n \times d}$$

denotes the design matrix.

The interpolation threshold occurs approximately when:

$$d \approx n.$$

At this point, the learner possesses just enough degrees of freedom to satisfy:

$$X\theta = y.$$

The interpolation manifold:

$$\mathcal{J} = \{\theta : X\theta = y\}$$

collapses toward low-dimensional or isolated solutions.

Near this threshold, small perturbations in data produce disproportionately large parameter changes.

Suppose:

$$\theta = (X^\top X)^{-1} X^\top y.$$

As:

$$\lambda_{\min}(X^\top X) \rightarrow 0,$$

the inverse becomes unstable:

$$\|(X^\top X)^{-1}\| \rightarrow \infty.$$

Variance therefore explodes near interpolation.

This produces the classical overfitting peak.

Within the representational regime framework, this corresponds to:

$$\rho \approx 1.$$

The learner possesses insufficient excess geometry to localize residual accom-

modations.

Local perturbations propagate globally.

B.3 The Double Descent Transition

Modern overparameterized systems violate the classical U-curve prediction.

As:

$$d \gg n,$$

generalization error frequently decreases again after the interpolation threshold.

This produces the double descent curve.

Formally:

$$\text{test error} = \mathcal{E}(d),$$

where:

$$\mathcal{E}(d)$$

initially decreases, then peaks near:

$$d \approx n,$$

and finally decreases again for:

$$d \gg n.$$

The second descent fundamentally contradicts classical assumptions that interpolation necessarily destroys generalization.

Within the regime framework, the explanation is geometric.

Once:

$$\rho \gg 1,$$

the learner regains sufficient representational freedom to isolate local accommodations into geometrically confined regions.

The system no longer requires global manifold distortion to satisfy interpolation constraints.

B.4 Linear Regression in the Overparameterized Regime

Consider linear regression:

$$y = X\theta + \epsilon.$$

Suppose:

$$d > n.$$

Then infinitely many interpolating solutions satisfy:

$$X\theta = y.$$

Gradient descent initialized at:

$$\theta_0 = 0$$

converges toward the minimum-norm interpolant:

$$\theta^* = X^\top (XX^\top)^{-1}y.$$

The resulting predictor satisfies:

$$\theta^* = \arg \min_{\theta} \{\|\theta\|_2 : X\theta = y\}.$$

This solution possesses remarkable properties.

Although interpolation is exact:

$$X\theta^* = y,$$

generalization may remain strong because:

$$\|\theta^*\|_2$$

is minimized globally.

Smoothness emerges through optimization geometry itself.

B.5 Benign Overfitting

Benign overfitting occurs when interpolation error vanishes:

$$\hat{f}(x_i) = y_i,$$

yet generalization error remains bounded:

$$\mathbb{E}[(\hat{f}(x) - f^*(x))^2] < \infty.$$

This phenomenon was historically considered impossible under classical learning theory.

Within the representational regime framework, however, benign overfitting becomes expected whenever:

$$\rho \gg 1,$$

and optimization dynamics preserve smooth manifold structure.

The learner absorbs noise locally while preserving global continuity.

B.6 Localized Accommodation Decomposition

Suppose:

$$\hat{f} = g + s,$$

where:

$$g$$

captures smooth latent structure, and:

$$s = \sum_i s_i$$

captures localized residual accommodation.

Each:

$$s_i$$

satisfies:

$$s_i(x_j) = 0 \quad \text{for } j \neq i,$$

approximately.

The interpolation problem becomes:

$$g(x_i) + s_i(x_i) = y_i.$$

Generalization remains strong whenever:

$$\|s_i(x)\| \rightarrow 0$$

rapidly away from:

$$x_i.$$

The learner therefore memorizes noise without globally destabilizing representation.

B.7 Kernel Regression and Infinite-Dimensional Interpolation

Kernel methods provide especially clear realizations of excess-capacity geometry.

Suppose:

$$K(x, x')$$

is a positive-definite kernel.

Kernel interpolation solves:

$$f(x) = \sum_{i=1}^n \alpha_i K(x, x_i),$$

subject to:

$$f(x_i) = y_i.$$

The reproducing kernel Hilbert space:

$$\mathcal{H}_K$$

may possess effectively infinite dimensionality.

Interpolation therefore occurs within extremely high-dimensional manifolds.

Smoothness emerges because kernel geometry constrains:

$$f$$

globally despite exact local interpolation.

This provides a concrete realization of:

$$g(x) + s_i(x)$$

decomposition.

B.8 Neural Tangent Kernels and Deep Networks

Overparameterized neural networks frequently behave approximately like kernel systems during optimization.

Suppose:

$$f_\theta(x)$$

denotes a neural network.

Linearizing near initialization:

$$\theta_0,$$

gives:

$$f_\theta(x) \approx f_{\theta_0}(x) + \nabla_\theta f_{\theta_0}(x)(\theta - \theta_0).$$

The induced neural tangent kernel is:

$$K_{\text{NTK}}(x, x') = \nabla_\theta f_{\theta_0}(x)^\top \nabla_\theta f_{\theta_0}(x').$$

Training therefore approximates kernel interpolation within high-dimensional representational geometry.

Generalization emerges through smoothness properties of the induced kernel manifold.

B.9 Flat Minima and Volume Concentration

One of the most important properties of high-dimensional optimization is volume concentration.

Suppose:

$$\mathcal{M}_{\text{flat}}$$

denotes low-curvature minima, and:

$$\mathcal{M}_{\text{sharp}}$$

denotes sharp minima.

In high-dimensional parameter spaces:

$$\text{Vol}(\mathcal{M}_{\text{flat}}) \gg \text{Vol}(\mathcal{M}_{\text{sharp}}).$$

Stochastic optimization therefore converges probabilistically toward flat minima.

This produces an implicit smoothness bias.

The learner prefers globally stable interpolation geometry because stable manifolds dominate parameter volume.

This provides a geometric explanation for soft inductive bias.

B.10 Grokking as Delayed Curvature Minimization

Grokking phenomena further support the framework.

Suppose:

$$\mathcal{L}_{\text{train}} \rightarrow 0,$$

while:

$$\mathcal{L}_{\text{test}}$$

remains initially high.

Extended optimization eventually produces:

$$\mathcal{L}_{\text{test}} \downarrow .$$

Within the representational regime framework, interpolation occurs first through fragmented local accommodations:

$$s_i(x).$$

Later optimization reorganizes global manifold smoothness:

$$g(x).$$

Generalization improves because:

$$g$$

stabilizes even though training interpolation was already complete.

Grokking therefore demonstrates that:

memorization

and:

manifold organization

occur on different optimization timescales.

B.11 Noise Scaling and Generalization Bounds

Suppose:

$$\epsilon_i$$

denotes observational noise.

Benign interpolation requires localized accommodation norms to remain bounded:

$$\sum_i \|s_i\|^2 < \infty.$$

As noise variance increases:

$$\sigma^2 \uparrow,$$

required accommodation curvature grows.

Eventually:

$$s_i$$

structures interfere globally.

This predicts a critical noise threshold:

$$\sigma_c^2,$$

beyond which interpolation becomes catastrophic.

The framework therefore predicts ecological dependence of learning regimes.

Excess capacity remains beneficial only when environmental structure permits localization of residual accommodation.

B.12 The Statistical Interpretation of Regime Dynamics

The representational regime framework may now be summarized statistically.

Constrained systems:

$$\rho \ll 1,$$

minimize variance by aggressively increasing bias.

Near-threshold systems:

$$\rho \approx 1,$$

exhibit unstable variance amplification.

Excess-capacity systems:

$$\rho \gg 1,$$

recover low variance through geometric localization of interpolation residuals.

Generalization therefore becomes compatible with interpolation whenever sufficient representational dimensionality exists to isolate local accommodations without globally destabilizing manifold smoothness.

This constitutes the formal statistical inversion underlying the entire framework.

The classical simplicity principle assumed:

$$\text{generalization} \Rightarrow \text{compression.}$$

Modern interpolation theory instead permits:

$$\text{generalization} \Rightarrow \text{organized abundance.}$$

C Dynamical Systems, Information Geometry, and Manifold Flows

The preceding appendices developed the representational regime framework primarily through the language of statistical learning theory and interpolation geometry. This appendix extends the formalism into dynamical systems theory, information geometry, and continuous manifold evolution.

The central objective is to reinterpret cognition not as static representation, but as evolving flow across structured representational state-spaces.

Under this perspective, learning becomes geometric motion.

C.1 Representational State-Spaces

Let:

$$\mathcal{Z}$$

denote the internal representational manifold of a learning system.

At any time:

$$t,$$

the learner occupies a representational state:

$$z(t) \in \mathcal{Z}.$$

Cognition therefore defines a trajectory:

$$\gamma : [0, T] \rightarrow \mathcal{Z},$$

with:

$$\gamma(t) = z(t).$$

The learner does not merely store representations statically.

The learner continuously traverses evolving representational geometry.

C.2 Learning as Gradient Flow

Suppose:

$$\mathcal{L}(z)$$

denotes a representational energy functional.

Learning dynamics may then be modeled as gradient flow:

$$\frac{dz}{dt} = -\nabla \mathcal{L}(z).$$

Under Euclidean geometry:

$$\nabla \mathcal{L}$$

corresponds to ordinary steepest descent.

However, cognition generally evolves on curved manifolds.

The appropriate dynamics therefore become:

$$\frac{dz^i}{dt} = -g^{ij} \frac{\partial \mathcal{L}}{\partial z^j},$$

where:

$$g^{ij}$$

denotes the inverse metric tensor on:

$$\mathcal{Z}.$$

Learning trajectories depend fundamentally upon representational geometry itself.

C.3 Information Geometry and Fisher Metrics

Suppose the learner represents probability distributions:

$$p_\theta(x).$$

Parameter space:

$$\Theta$$

acquires natural Riemannian structure through the Fisher information metric:

$$F_{ij} = \mathbb{E} \left[\frac{\partial \log p_\theta}{\partial \theta_i} \frac{\partial \log p_\theta}{\partial \theta_j} \right].$$

The resulting manifold:

$$(\Theta, F)$$

defines the intrinsic geometry of representational distinguishability.

Distances correspond not merely to parameter differences, but to informational discrimination capacity.

The representational regime framework therefore acquires a geometric reinterpretation:

$$C$$

corresponds approximately to effective navigable curvature volume within:

$$(\Theta, F).$$

High-capacity systems possess large navigable informational manifolds.

Constrained systems occupy highly compressed informational geometries.

C.4 Geodesic Interpolation

Suppose experiences:

$$x_i$$

and:

$$x_j$$

occupy representational states:

$$z_i, \quad z_j.$$

Interpolation corresponds to constructing geodesic trajectories:

$$\gamma_{ij}(t)$$

minimizing:

$$\ell(\gamma) = \int_0^1 \sqrt{g_{ab} \frac{d\gamma^a}{dt} \frac{d\gamma^b}{dt}} dt.$$

Generalization emerges when unseen experiences lie near low-curvature geodesic neighborhoods connecting prior experiences.

The learner predicts effectively because representational manifolds remain smooth.

This geometric interpretation directly connects:

$$g(x)$$

global continuity with geodesic stability.

C.5 Curvature and Cognitive Rigidity

Representational curvature determines cognitive flexibility.

Suppose:

$$R_{ijkl}$$

denotes the Riemann curvature tensor on:

$$\mathcal{Z}.$$

Highly curved regions correspond to brittle local attractors.

Small perturbations produce large trajectory divergence:

$$\delta z(t) \uparrow .$$

Near the interpolation threshold:

$$\rho \approx 1,$$

representational curvature becomes unstable.

The learner develops rigid high-curvature interpolation structures.

This predicts brittle cognition, excessive sensitivity to local perturbation, and poor transfer.

Excess-capacity systems instead distribute curvature more smoothly across high-dimensional manifolds:

$$|R_{ijkl}| \downarrow .$$

Global stability increases.

C.6 Attractor Dynamics

Learning systems frequently exhibit attractor structure.

Suppose:

$$A \subset \mathcal{Z}$$

denotes an attractor manifold satisfying:

$$\lim_{t \rightarrow \infty} d(z(t), A) = 0.$$

Concepts, habits, perceptual categories, and cognitive routines may all correspond to attractor geometries.

Constraint therefore becomes topological stabilization.

Healthy cognition requires:

stable attractors

without:

rigid trapping.

Excessive rigidity produces pathological fixation:

$$z(t) \rightarrow A_i$$

for narrow:

$$A_i.$$

Insufficient stabilization produces semantic diffusion:

$$z(t)$$

wanders chaotically through:

$$\mathcal{Z}.$$

Healthy cognition occupies metastable manifold organization.

C.7 Entropy and Trajectory Volume

Representational abundance may also be characterized dynamically through trajectory entropy.

Suppose:

$$\mathcal{T}(z_0)$$

denotes the set of admissible future trajectories from initial state:

$$z_0.$$

Define trajectory entropy:

$$S(z_0) = \log \text{Vol}(\mathcal{T}(z_0)).$$

High trajectory entropy corresponds to flexible representational exploration.

Low trajectory entropy corresponds to rigid constrained cognition.

The representational regime framework predicts:

$$\rho \uparrow \Rightarrow S \uparrow.$$

Excess-capacity systems preserve larger volumes of admissible manifold traversal.

However, excessive entropy without stabilization produces navigational chaos.

Constraint therefore functions by regulating trajectory entropy.

C.8 Attention as Curvature Navigation

Attention may be formalized geometrically.

Suppose:

$$\mathcal{A}(t)$$

denotes the attentional trajectory over:

$$\mathcal{Z}.$$

Attention follows informational gradients:

$$\nabla\Phi(z),$$

where:

$$\Phi$$

denotes representational salience.

The dynamics become:

$$\frac{d\mathcal{A}}{dt} = \nabla\Phi(z).$$

Flat manifolds:

$$\nabla\Phi \approx 0$$

produce boredom and disengagement.

Excessively fragmented manifolds:

$$\nabla\Phi$$

chaotic, produce attentional instability.

Healthy cognition requires smoothly structured informational gradients.

This directly connects exploratory engagement with manifold curvature organization.

C.9 Prediction Error and Manifold Correction

Predictive systems continuously compare expected trajectories:

$$\hat{z}(t)$$

against observed trajectories:

$$z(t).$$

Prediction error becomes:

$$\epsilon(t) = z(t) - \hat{z}(t).$$

Learning updates manifold geometry to minimize:

$$\mathcal{E} = \int \|\epsilon(t)\|^2 dt.$$

Constrained systems minimize:

$$\mathcal{E}$$

through compression.

Excess-capacity systems minimize:

$$\mathcal{E}$$

through localized accommodation.

The distinction corresponds to different geometric correction strategies.

C.10 Catastrophic Overfitting as Curvature Explosion

Suppose local accommodations:

$$s_i(x)$$

begin interacting globally.

Then manifold curvature increases rapidly:

$$|R_{ijkl}| \rightarrow \infty.$$

Nearby trajectories diverge chaotically.

Generalization collapses.

Catastrophic overfitting therefore corresponds geometrically to curvature explosion.

Benign overfitting remains possible only when localized accommodations preserve bounded curvature almost everywhere.

C.11 Neural Collapse and Symmetry Formation

Recent work on neural collapse suggests that trained deep systems frequently converge toward symmetric representational configurations.

Class means become approximately equiangular:

$$\langle \mu_i, \mu_j \rangle \approx -\frac{1}{K-1},$$

for:

$$i \neq j.$$

This implies emergent simplex geometry within representation space.

The representational regime framework interprets neural collapse as spontaneous manifold regularization.

High-dimensional interpolation manifolds reorganize toward maximally symmetric low-curvature global configurations.

Generalization emerges because representation geometry self-organizes into stable topological partitions.

C.12 Representational Phase Transitions

The framework predicts genuine representational phase transitions.

Suppose:

$$\rho = \frac{C}{E}.$$

Critical values:

$$\rho_c$$

separate qualitatively distinct manifold geometries.

As:

$$\rho$$

crosses:

$$\rho_c,$$

the topology of admissible interpolation manifolds changes abruptly.

Possible transitions include:

fragmented \rightarrow globally connected,

rigid \rightarrow smooth,

compressed \rightarrow abundant,

or:

symbolic \rightarrow distributed.

Scaling laws in artificial intelligence may therefore reflect geometric phase transitions rather than continuous accumulation alone.

C.13 The Dynamical Interpretation of Intelligence

The dynamical framework developed throughout this appendix permits a final reinterpretation of intelligence itself.

Intelligence is not merely:

compression,

nor:

optimization,

nor:

prediction.

Intelligence becomes stable manifold traversal within high-dimensional representational geometry.

The learner succeeds when trajectories remain globally coherent despite dense local accommodation.

Healthy cognition preserves:

high trajectory flexibility,

together with:

low global instability.

This balance defines the geometry of adaptive cognition.

The mind therefore becomes neither a symbolic machine nor a statistical table.

It becomes a dynamical manifold navigation system operating under conditions of representational abundance, environmental constraint, and continuous geometric reorganization.

D Functional Analysis, Kernel Operators, and Spectral Geometry

This appendix develops a deeper functional-analytic formulation of the representational regime framework. The previous appendices treated cognition geometrically and statistically. Here, the framework is reformulated in terms of operators, reproducing kernels, spectral decompositions, and function spaces.

The central objective is to show that excess-capacity learning may be understood as a problem of spectral organization across high-dimensional representational operators.

D.1 Representational Operators

Let:

$$\mathcal{X}$$

denote experiential space and:

$$\mathcal{H}$$

a Hilbert space of admissible representations.

A cognitive system defines an operator:

$$T : \mathcal{H} \rightarrow \mathcal{H},$$

mapping representational states into predictive transformations.

Experiences:

$$(x_i, y_i)$$

impose interpolation constraints:

$$Tf(x_i) = y_i.$$

The learner therefore constructs:

$$f \in \mathcal{H}$$

satisfying operator constraints induced by environmental interaction.

D.2 Kernel Representations

Suppose:

$$K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$$

is a positive-definite kernel.

Then:

$$K$$

induces a reproducing kernel Hilbert space:

$$\mathcal{H}_K.$$

Every:

$$f \in \mathcal{H}_K$$

satisfies:

$$f(x) = \langle f, K_x \rangle_{\mathcal{H}_K},$$

where:

$$K_x(\cdot) = K(x, \cdot).$$

Representations therefore become superpositions of localized similarity operators:

$$f(x) = \sum_{i=1}^n \alpha_i K(x, x_i).$$

The representational regime framework interprets cognition as the organization of such localized operators within globally coherent manifolds.

D.3 Localized Accommodation as Spectral Perturbation

Suppose:

$$f = g + s,$$

where:

$$g$$

denotes smooth global structure, and:

$$s = \sum_i s_i$$

localized accommodation operators.

Each:

$$s_i$$

acts approximately as a localized perturbation:

$$s_i(x) \approx \alpha_i K(x, x_i).$$

The learner therefore constructs:

$$f$$

through spectral superposition of globally smooth low-frequency structure together with highly localized high-frequency corrections.

This decomposition parallels Fourier analysis:

$$f(x) = \sum_k \hat{f}_k e^{ikx}.$$

Global smoothness corresponds to dominance of low-frequency spectral modes.

Localized spikes correspond to confined high-frequency spectral energy.

D.4 Spectral Bias

Modern neural systems exhibit spectral bias: low-frequency structure is learned before high-frequency detail.

Suppose:

$$f^* = \sum_k \hat{f}_k \phi_k,$$

where:

$$\phi_k$$

are eigenfunctions of an operator:

$$L.$$

Gradient-based learning frequently converges first toward:

$$|\hat{f}_k|$$

with small:

$$\lambda_k,$$

where:

$$\lambda_k$$

denotes spectral frequency.

The learner therefore first constructs:

$$g(x),$$

global manifold continuity, before gradually incorporating:

$$s_i(x),$$

localized residual accommodations.

This provides a spectral interpretation of grokking and delayed generalization.

D.5 Operator Smoothness and Sobolev Norms

Global representational smoothness may be formalized through Sobolev functionals.

Suppose:

$$f \in H^m(\mathcal{X}),$$

a Sobolev space of order:

$$m.$$

Define:

$$\|f\|_{H^m}^2 = \sum_{|\alpha| \leq m} \int |D^\alpha f(x)|^2 dx.$$

Smooth interpolation minimizes high-order derivatives globally while preserving exact local constraints.

Excess-capacity systems therefore solve:

$$\min_f \|f\|_{H^m}$$

subject to:

$$f(x_i) = y_i.$$

This formalizes the claim that generalization emerges through smooth interpolation geometry rather than compression alone.

D.6 Mercer Decomposition and Cognitive Modes

By Mercer's theorem:

$$K(x, y) = \sum_{k=1}^{\infty} \lambda_k \phi_k(x) \phi_k(y).$$

Representational geometry therefore decomposes into orthogonal cognitive modes:

$$\phi_k.$$

Constrained systems preserve only dominant low-frequency modes:

$$\lambda_1, \lambda_2, \dots$$

Excess-capacity systems preserve substantially richer spectral structure:

$$\lambda_k$$

across many frequencies.

This predicts that high-capacity cognition retains detailed local structure because higher-frequency representational modes remain accessible.

D.7 Compression as Spectral Truncation

The classical simplicity principle may now be reformulated spectrally.

Compression-centered cognition corresponds approximately to spectral truncation:

$$f_N(x) = \sum_{k=1}^N \lambda_k \phi_k(x),$$

with:

$$N$$

small.

The learner discards higher-frequency representational structure:

$$k > N.$$

This produces stable low-complexity generalization but eliminates local accommodation capacity.

Excess-capacity learning instead preserves broader spectral bandwidth.

Generalization emerges not through eliminating high-frequency structure entirely, but through localizing it spectrally.

D.8 The Spectral Interpretation of Overfitting

Classical overfitting corresponds to uncontrolled amplification of high-frequency modes.

Suppose:

$$|\hat{f}_k| \rightarrow \infty$$

for large:

$$k.$$

The representational manifold becomes highly oscillatory.

Generalization collapses.

Benign overfitting differs critically.

Localized accommodations remain spectrally confined:

$$\sum_{k>K} |\hat{f}_k|^2$$

bounded.

The learner preserves low-frequency global continuity while isolating high-frequency corrections into narrow localized regions.

This provides a spectral explanation for:

$$g(x) + s_i(x)$$

decomposition.

D.9 Attention as Operator Reweighting

Attention mechanisms may also be formulated spectrally.

Suppose:

$$A : \mathcal{H} \rightarrow \mathcal{H}$$

denotes an attentional operator.

Attention dynamically reweights spectral components:

$$Af = \sum_k a_k \hat{f}_k \phi_k.$$

Focused attention amplifies selected representational frequencies while suppressing others.

This predicts that cognition dynamically reorganizes spectral geometry depending upon task structure and environmental demands.

Attention becomes adaptive spectral routing.

D.10 Memory as Operator Persistence

Memory corresponds to persistence of representational operators across time.

Suppose:

$$T_t$$

denotes the representational operator at time:

$$t.$$

Memory stability requires:

$$\|T_{t+\Delta t} - T_t\|$$

small for relevant experiential structures.

Catastrophic forgetting occurs when:

$$\|T_{t+\Delta t} - T_t\|$$

large.

Excess-capacity systems reduce catastrophic forgetting by embedding memories into geometrically separated operator regions.

Interference decreases because spectral overlap weakens.

D.11 The Spectral Geometry of Expertise

Expertise corresponds to spectral refinement.

Novice systems rely primarily upon low-frequency coarse structure:

$$\lambda_1, \lambda_2.$$

Experts preserve richer localized spectral detail:

$$\lambda_k$$

for large:

$$k.$$

However, healthy expertise maintains global low-frequency continuity simultaneously.

Expert cognition therefore corresponds to multiscale spectral organization.

The learner preserves:

global smoothness

and:

localized refinement.

This directly parallels the excess-capacity framework developed throughout the essay.

D.12 Operator Phase Transitions

Representational regimes may also be characterized spectrally.

Constrained systems exhibit strong spectral concentration:

$$\sum_{k>N} \lambda_k \approx 0.$$

Near-threshold systems exhibit unstable high-frequency amplification:

$$\lambda_k \rightarrow \infty$$

irregularly.

Excess-capacity systems exhibit distributed but organized spectral richness:

$$\lambda_k$$

broadly supported yet geometrically localized.

Phase transitions therefore correspond to changes in spectral organization across representational operators.

D.13 Toward a Spectral Theory of Cognition

The broader implication of this appendix is that cognition may ultimately require a fundamentally spectral interpretation.

Learning systems organize information not merely symbolically, but through dynamic spectral geometry distributed across high-dimensional operator manifolds.

Compression becomes spectral truncation.

Interpolation becomes spectral superposition.

Generalization becomes low-curvature spectral continuity.

Memory becomes operator persistence.

Attention becomes adaptive spectral routing.

Expertise becomes multiscale spectral organization.

The mind therefore appears less like a symbolic compression engine and more like a continuously evolving spectral manifold dynamically balancing localized accommodation against global continuity.

The representational abundance framework thus converges naturally with functional analysis, kernel methods, harmonic analysis, and spectral geometry.

Intelligence becomes the structured organization of representational spectra across dynamically evolving manifolds.

E Category Theory, Sheaf Structures, and Cognitive Gluing

The previous appendices developed the representational regime framework through geometry, statistics, dynamics, and spectral analysis. This appendix extends the

framework into category theory and sheaf-theoretic organization.

The objective is to formalize cognition as the structured gluing of local representational neighborhoods into globally coherent manifold structure.

Under this interpretation, intelligence becomes a problem of coherent compositionality across distributed local representations.

E.1 Representations as Morphisms

Let:

$$\mathcal{C}$$

be a category whose objects represent experiential states and whose morphisms represent admissible cognitive transformations.

For:

$$X, Y \in \text{Ob}(\mathcal{C}),$$

a morphism:

$$f : X \rightarrow Y$$

represents a learned transition, inference, abstraction, or predictive transformation.

Cognition therefore becomes compositional:

$$g \circ f : X \rightarrow Z.$$

The learner organizes not isolated symbols, but networks of transformable relationships.

E.2 Representational Functors

Suppose:

$$\mathcal{E}$$

denotes the category of experiential situations, and:

$$\mathcal{R}$$

the category of internal representations.

A cognitive system defines a functor:

$$F : \mathcal{E} \rightarrow \mathcal{R}.$$

The functor preserves structural relationships:

$$F(g \circ f) = F(g) \circ F(f).$$

Learning therefore corresponds to constructing structure-preserving mappings between experiential geometry and representational organization.

Compression-centered theories implicitly assume that:

$$F$$

acts primarily through reduction.

The representational abundance framework instead permits:

$$F$$

to increase representational dimensionality while preserving relational structure.

E.3 Local Representational Neighborhoods

Suppose:

$$\mathcal{U} = \{U_i\}$$

is an open cover over experiential manifold:

$$\mathcal{X}.$$

Each:

$$U_i$$

corresponds to a localized experiential neighborhood.

The learner constructs local representational sections:

$$s_i \in \mathcal{F}(U_i),$$

where:

$$\mathcal{F}$$

is a representational presheaf.

Each:

$$s_i$$

encodes local interpolation structure.

Localized accommodations:

$$s_i(x)$$

therefore acquire a sheaf-theoretic interpretation.

E.4 Gluing Conditions

A central problem of cognition is global coherence.

Local representational sections must satisfy compatibility conditions on overlaps:

$$U_i \cap U_j.$$

Specifically:

$$s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}.$$

When compatible local sections glue successfully, there exists a global section:

$$s \in \mathcal{F}(\mathcal{X}),$$

satisfying:

$$s|_{U_i} = s_i.$$

Generalization therefore corresponds to successful gluing of local experiential neighborhoods into globally coherent representational structure.

E.5 Compression as Forced Globalization

Classical simplicity-centered theories aggressively force local sections into low-dimensional global summaries.

The learner minimizes local differentiation:

$$s_i \approx s_j$$

globally.

This produces stable but lossy abstraction.

The representational abundance framework instead permits richer local structure while maintaining compatibility on overlaps.

Generalization emerges through coherent gluing rather than forced homogenization.

E.6 Localized Accommodation and Cohomological Obstruction

Not all local representations glue globally.

Suppose:

$$s_i$$

and:

$$s_j$$

disagree on:

$$U_i \cap U_j.$$

The mismatch defines a cocycle:

$$c_{ij} = s_i - s_j.$$

Global inconsistency corresponds to nontrivial cohomology:

$$H^1(\mathcal{X}, \mathcal{F}) \neq 0.$$

Cognitive contradiction, ambiguity, or unresolved conceptual tension may therefore correspond to cohomological obstruction.

The learner fails to globally integrate local representational neighborhoods.

E.7 Excess Capacity and Local Resolution

Excess-capacity systems possess additional representational freedom permitting localized accommodation without globally destabilizing representation.

Formally, high-dimensional representational geometry enlarges the space of

admissible local sections:

$$\mathcal{F}(U_i) \uparrow .$$

The learner therefore gains flexibility to resolve local inconsistencies through additional representational coordinates.

Instead of collapsing incompatible experiences through aggressive compression, the learner preserves differentiated local structure while maintaining global compatibility through richer gluing geometry.

This formalizes the claim that excess-capacity systems preserve experiential richness without sacrificing global coherence.

E.8 Representational Fiber Bundles

The framework may also be formulated through fiber bundle structure.

Suppose:

$$\pi : E \rightarrow \mathcal{X}$$

is a fiber bundle.

For each:

$$x \in \mathcal{X},$$

the fiber:

$$\pi^{-1}(x)$$

represents the space of admissible interpretations, predictions, or local representational embeddings associated with:

$$x.$$

Learning corresponds to constructing smooth sections:

$$s : \mathcal{X} \rightarrow E.$$

Constrained systems possess highly restricted fibers:

$$\dim(\pi^{-1}(x))$$

small.

Excess-capacity systems possess large fibers supporting richer local accommodation.

Generalization depends upon constructing globally smooth sections across these fibers.

E.9 Attention as Local Trivialization

Attention may also be interpreted categorically.

Suppose:

$$\mathcal{A}$$

selects local coordinate charts:

$$(U_i, \phi_i).$$

Attention dynamically trivializes local representational neighborhoods:

$$\pi^{-1}(U_i) \cong U_i \times F.$$

The learner temporarily simplifies local geometry to facilitate efficient interpolation and inference.

Attention therefore functions as dynamic local coordinate selection within high-dimensional representational bundles.

E.10 Functorial Generalization

Generalization itself may be interpreted functorially.

Suppose:

$$F : \mathcal{E} \rightarrow \mathcal{R}$$

preserves local relational structure.

Then unseen experiences:

$$x^*$$

can be integrated coherently whenever:

$$F$$

extends smoothly over neighboring representational regions.

Generalization therefore becomes extension of functorial coherence across manifold neighborhoods.

This directly parallels interpolation geometry developed throughout the main text.

E.11 Natural Transformations and Learning Dynamics

Suppose:

$$F_t$$

denotes the representational functor at time:

$$t.$$

Learning corresponds to sequences of natural transformations:

$$\eta_t : F_t \Rightarrow F_{t+1}.$$

Cognitive development therefore becomes deformation of representational functors through time.

Healthy learning preserves sufficient structural continuity:

$$\eta_t$$

smooth, while permitting local representational refinement.

Catastrophic forgetting corresponds to discontinuous functor deformation.

E.12 Hierarchical Sheaf Structure

Cognition exhibits strong multiscale organization.

Local perceptual neighborhoods glue into intermediate conceptual structures, which themselves glue into broader semantic manifolds.

This naturally suggests hierarchical sheaf structure:

$$\mathcal{F}_1 \subset \mathcal{F}_2 \subset \dots \subset \mathcal{F}_n.$$

Lower levels preserve local sensory differentiation.

Higher levels preserve broad semantic continuity.

Excess-capacity learning permits coexistence of both scales simultaneously.

The learner avoids premature collapse of local structure into rigid global abstraction.

E.13 The One-System Hypothesis Revisited Categorically

The category-theoretic framework provides additional support for the moderate one-system hypothesis developed earlier.

Different cognitive functions need not correspond to fundamentally separate systems.

Instead, apparently distinct cognitive behaviors may correspond to different regions, scales, or gluing structures within unified representational categories.

Episodic memory corresponds to localized sections.

Semantic abstraction corresponds to globally coherent sheaf structure.

Attention corresponds to local trivialization.

Inference corresponds to morphism composition.

Learning corresponds to functor deformation.

The same representational substrate supports all of these processes through different organizational geometries.

E.14 Toward a Sheaf-Theoretic Theory of Cognition

The broader implication of this appendix is that cognition may fundamentally require sheaf-theoretic organization.

The learner continuously constructs local representations, checks overlap compatibility, resolves gluing inconsistencies, and stabilizes globally coherent manifold structure.

Compression becomes forced gluing under insufficient representational freedom.

Excess capacity permits richer local differentiation while preserving compatibility globally.

Generalization emerges not because local detail disappears, but because local detail becomes coherently organized across representational neighborhoods.

The mind therefore appears not merely as a symbolic machine nor solely as a statistical interpolator.

It becomes a distributed gluing system operating across dynamically evolving representational manifolds.

F Thermodynamics, Entropy, and Representational Phase Space

The preceding appendices developed the representational regime framework through statistical learning theory, geometry, dynamics, spectral analysis, and categorical organization. This appendix introduces a thermodynamic interpretation of cognition.

The objective is to formalize learning systems as entropy-regulating dynamical structures evolving through representational phase spaces under energetic, informational, and geometric constraints.

Under this interpretation, cognition becomes a problem of managing representational entropy while preserving manifold stability.

F.1 Representational Phase Spaces

Let:

$$\Gamma$$

denote the representational phase space of a cognitive system.

Each state:

$$z \in \Gamma$$

corresponds to a complete representational configuration.

The system evolves dynamically:

$$z(t) : [0, T] \rightarrow \Gamma.$$

Learning therefore corresponds to trajectories through representational phase space.

Unlike classical thermodynamic systems, however, cognitive systems actively reshape their own phase geometry during evolution.

The manifold itself changes through learning.

F.2 Representational Entropy

Suppose:

$$p(z)$$

denotes the distribution over representational states.

Define representational entropy:

$$S = - \int_{\Gamma} p(z) \log p(z) dz.$$

This entropy measures uncertainty over representational configurations.

High entropy corresponds to diffuse manifold organization.

Low entropy corresponds to rigid highly constrained structure.

Healthy cognition generally occupies intermediate entropy regimes:

$$0 < S < \infty.$$

Excessively low entropy produces rigidity and premature compression.

Excessively high entropy produces incoherent semantic diffusion.

F.3 Experiential Entropy and Capacity

Suppose experiential complexity:

$$E$$

is measured through trajectory entropy:

$$E = \log \text{Vol}(\mathcal{T}),$$

where:

$$\mathcal{T}$$

denotes admissible experiential trajectories.

Representational flexibility:

$$C$$

determines how many distinct trajectories can be stably embedded into representational geometry.

The operational regime ratio:

$$\rho = \frac{C}{E}$$

therefore acquires a thermodynamic interpretation.

Constrained systems:

$$\rho \ll 1,$$

lack sufficient representational phase volume.

Experiential trajectories collide.

Compression becomes necessary.

Excess-capacity systems:

$$\rho \gg 1,$$

possess sufficiently large phase spaces to preserve substantial local differentiation.

F.4 Free Energy and Predictive Stabilization

Suppose the learner maintains internal predictions:

$$q(z),$$

while the environment generates actual observations:

$$p(z).$$

Define variational free energy:

$$\mathcal{F} = \mathbb{E}_q[\log q(z) - \log p(z, x)].$$

Minimizing:

$$\mathcal{F}$$

reduces prediction error while preserving representational stability.

Within the representational regime framework, constrained and excess-capacity systems minimize free energy differently.

Constrained systems reduce:

$$\mathcal{F}$$

through aggressive compression:

$$q(z)$$

low-dimensional.

Excess-capacity systems instead preserve richer representational distributions while stabilizing prediction through smooth manifold organization.

The distinction becomes:

compression-driven minimization

versus:

interpolation-driven minimization.

F.5 Representational Temperature

Introduce a representational temperature parameter:

$$T.$$

High:

$$T$$

corresponds to exploratory flexible manifold traversal.

Low:

$$T$$

corresponds to rigid attractor stabilization.

The learner's representational distribution becomes:

$$p(z) = \frac{1}{Z} e^{-E(z)/T},$$

where:

$$E(z)$$

is representational energy and:

$$Z$$

the partition function:

$$Z = \int_{\Gamma} e^{-E(z)/T} dz.$$

Development, learning, exploration, and pathology may all be interpreted through changes in representational temperature.

High-temperature cognition supports broad interpolation and creative recombination.

Low-temperature cognition supports stable precision and efficient execution.

Healthy systems dynamically regulate:

$$T.$$

F.6 Phase Transitions Between Regimes

The framework predicts genuine thermodynamic phase transitions.

Suppose:

$$\rho = \frac{C}{E}.$$

Critical values:

$$\rho_c$$

separate qualitatively distinct organizational phases.

As:

$$\rho$$

crosses:

$$\rho_c,$$

the topology of representational phase space reorganizes abruptly.

Possible transitions include:

compressed \rightarrow distributed,

rigid \rightarrow fluid,

fragmented \rightarrow globally coherent.

These transitions resemble symmetry-breaking phenomena in statistical physics.

Scaling behavior in neural systems and artificial intelligence may therefore reflect thermodynamic reorganization rather than continuous incremental accumulation alone.

F.7 Localized Accommodation as Entropic Isolation

Localized accommodation structures:

$$s_i(x)$$

may be interpreted thermodynamically as entropic isolation mechanisms.

Excess-capacity systems preserve local irregularities by confining them into low-volume regions of representational phase space.

Noise becomes trapped within localized energetic basins.

Global manifold organization remains stable because entropy generated by local irregularity does not propagate globally.

This provides a thermodynamic interpretation of benign overfitting.

F.8 Catastrophic Overfitting as Entropic Propagation

Suppose localized accommodations fail to remain isolated.

Then local perturbations propagate across:

$$\Gamma.$$

Representational entropy increases globally:

$$S \uparrow .$$

The manifold destabilizes.

Prediction error amplifies.

Catastrophic overfitting therefore corresponds thermodynamically to uncon-

trolled entropy propagation across representational phase space.

Benign overfitting remains possible only when local entropy production remains geometrically confined.

F.9 Memory as Entropy Stabilization

Memory may also be interpreted thermodynamically.

Stable memories correspond to low-energy attractor basins:

$$A_i \subset \Gamma.$$

A memory persists when:

$$\Delta E(A_i)$$

remains sufficiently large relative to thermal fluctuations:

$$T.$$

Catastrophic forgetting occurs when attractor barriers weaken:

$$\Delta E(A_i) \downarrow.$$

Excess-capacity systems preserve memory by constructing geometrically separated attractor basins with reduced interference.

F.10 Attention and Entropic Gradients

Attention corresponds to movement along representational free-energy gradients.

Suppose:

$$\Phi(z)$$

denotes attentional potential.

Attention dynamics satisfy:

$$\frac{dz}{dt} = -\nabla\Phi(z).$$

Regions with high informational curvature attract representational flow.

Boredom corresponds to flattened entropic landscapes:

$$\nabla\Phi \approx 0.$$

Anxiety corresponds to fragmented unstable landscapes:

$$\nabla\Phi$$

chaotic.

Healthy attention requires structured intermediate entropy gradients supporting stable exploratory traversal.

F.11 Creativity and Metastable Exploration

Creativity emerges naturally within metastable representational systems.

Suppose:

$$A_1, A_2, \dots, A_n$$

denote partially stable attractor basins.

Creative cognition corresponds to transitions between attractors without complete collapse into random diffusion.

The learner traverses:

$$A_i \rightarrow A_j$$

through intermediate metastable trajectories.

Excess-capacity systems support creativity because representational abundance preserves large admissible trajectory volumes:

$$\text{Vol}(\mathcal{T}) \uparrow.$$

Constraint remains necessary to prevent complete thermodynamic diffusion.

Creativity therefore emerges from constrained metastability.

F.12 Biological Scaling and Energetic Tradeoffs

The thermodynamic interpretation also clarifies biological tradeoffs.

Increasing representational abundance requires increased energetic mainte-

nance:

$$\mathcal{E}_{\text{metabolic}} \uparrow .$$

High-dimensional representational manifolds are metabolically expensive.

However, abundance also increases flexibility, robustness, interpolation fidelity, and adaptive trajectory volume.

Biological intelligence therefore appears to balance:

energetic minimization,

against:

representational flexibility.

Evolution may thus optimize not simplicity alone, but thermodynamically efficient manifold organization.

F.13 Toward a Thermodynamic Theory of Cognition

The broader implication of this appendix is that cognition may fundamentally require thermodynamic interpretation.

Learning systems continuously regulate entropy, stabilize attractor geometries, localize perturbations, preserve manifold continuity, and balance exploration against energetic cost.

Compression becomes entropy reduction through representational collapse.

Interpolation becomes entropy organization through geometric isolation.

Generalization becomes thermodynamic stabilization across high-dimensional manifold structure.

The mind therefore appears not merely computational, symbolic, or statistical.

It becomes a thermodynamic manifold-regulation system operating within evolving representational phase spaces.

G Variational Principles and the Euler–Lagrange Formulation of Representational Dynamics

The previous appendices developed the representational regime framework through geometry, statistical learning theory, spectral analysis, category theory, and thermodynamics. This appendix introduces a variational formulation unifying these perspectives.

The central objective is to derive cognition from an action principle.

Under this interpretation, learning systems evolve by extremizing representational functionals balancing interpolation fidelity, manifold smoothness, energetic efficiency, and entropy regulation.

Cognition becomes variational geometry.

G.1 Representational Action Functionals

Let:

$$\mathcal{Z}$$

denote the representational manifold.

A cognitive trajectory:

$$z(t)$$

defines a path:

$$\gamma : [0, T] \rightarrow \mathcal{Z}.$$

We define the representational action:

$$\mathcal{A}[z] = \int_0^T \mathcal{L}(z, \dot{z}, t) dt,$$

where:

$$\mathcal{L}$$

is a representational Lagrangian.

Learning corresponds to trajectories extremizing:

$$\delta \mathcal{A} = 0.$$

G.2 The Representational Lagrangian

We decompose:

$$\mathcal{L} = \mathcal{K} - \mathcal{V} - \mathcal{R},$$

where:

$$\mathcal{K}$$

represents representational flexibility,

$$\mathcal{V}$$

represents prediction error or constraint mismatch,

and:

$$\mathcal{R}$$

represents regularization or curvature cost.

A canonical form is:

$$\mathcal{L} = \frac{1}{2}g_{ij}(z)\dot{z}^i\dot{z}^j - \Phi(z) - \lambda\|\nabla^2 z\|^2.$$

The first term measures representational motion across manifold geometry.

The second term measures predictive inconsistency.

The third term penalizes excessive curvature.

G.3 Euler–Lagrange Dynamics

The representational trajectory satisfies:

$$\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial \dot{z}^i} \right) - \frac{\partial \mathcal{L}}{\partial z^i} = 0.$$

Substituting the representational Lagrangian gives:

$$\frac{d}{dt} (g_{ij}\dot{z}^j) - \frac{1}{2} \frac{\partial g_{jk}}{\partial z^i} \dot{z}^j \dot{z}^k + \frac{\partial \Phi}{\partial z^i} + \lambda \frac{\partial}{\partial z^i} \|\nabla^2 z\|^2 = 0.$$

Cognition therefore evolves through competition between:

manifold traversal,
prediction stabilization,

and:

curvature minimization.

G.4 Interpolation Constraints as Boundary Conditions

Experiences:

$$(x_i, y_i)$$

impose interpolation constraints.

Suppose:

$$z(t_i) = y_i.$$

The variational problem becomes:

$$\delta\mathcal{A}[z] = 0$$

subject to:

$$z(t_i) = y_i.$$

Constrained systems lack sufficient representational freedom to satisfy all constraints simultaneously while preserving smoothness.

Excess-capacity systems possess sufficient degrees of freedom to satisfy interpolation locally while minimizing curvature globally.

This formalizes the:

$$g(x) + s_i(x)$$

decomposition variationally.

G.5 Localized Accommodation Fields

Suppose:

$$z(t) = g(t) + s(t),$$

where:

$$g$$

is globally smooth manifold structure, and:

$$s = \sum_i s_i$$

localized accommodation fields.

The action becomes:

$$\mathcal{A} = \mathcal{A}_g + \mathcal{A}_s + \mathcal{A}_{\text{int}},$$

where:

$$\mathcal{A}_{\text{int}}$$

captures coupling between global smoothness and local accommodation.

Benign overfitting occurs when:

$$\mathcal{A}_{\text{int}}$$

remains bounded.

Catastrophic overfitting occurs when local accommodation fields destabilize global manifold curvature.

G.6 Representational Curvature Minimization

The framework predicts that learning systems minimize curvature functionals globally whenever representational abundance permits.

Define:

$$\mathcal{R}[z] = \int \|\nabla^2 z\|^2 dx.$$

Then excess-capacity learning solves:

$$\min_z \mathcal{R}[z]$$

subject to:

$$z(x_i) = y_i.$$

The resulting Euler–Lagrange equation becomes:

$$\Delta^2 z = 0$$

almost everywhere away from localized interpolation constraints.

Thus globally smooth interpolation emerges naturally.

G.7 Noether Symmetries and Cognitive Invariants

Variational systems possess conserved quantities under continuous symmetry transformations.

Suppose:

$$\mathcal{L}$$

is invariant under:

$$z \rightarrow z + \epsilon\eta.$$

Then Noether's theorem implies conservation laws.

Potential cognitive invariants include:

semantic continuity,

trajectory coherence,

predictive stability,

or:

attentional persistence.

Healthy cognition therefore preserves invariant manifold structure despite continual local adaptation.

G.8 Attention as Variational Focusing

Attention may also be derived variationally.

Suppose attentional allocation:

$$a(x)$$

weights representational energy.

The modified action becomes:

$$\mathcal{A} = \int a(x)\mathcal{L}(x)dx.$$

Attention therefore acts by dynamically reshaping representational curvature landscapes.

Focused attention amplifies selected manifold regions while suppressing irrelevant directions.

This provides a variational interpretation of adaptive salience.

G.9 Learning as Geodesic Flow

Suppose:

$$\Phi(z) = 0.$$

Then cognition reduces to geodesic flow:

$$\frac{D\dot{z}}{dt} = 0.$$

Learning follows minimal-action trajectories through representational geometry.

Prediction error introduces external forcing terms:

$$\nabla\Phi(z).$$

The learner continuously adjusts geodesic structure to maintain interpolation consistency.

G.10 Constraint Potentials and Cognitive Rigidity

Constraint may be represented through potential wells:

$$\Phi(z).$$

Deep narrow wells produce rigid attractor dynamics.

Shallow broad wells produce flexible metastable organization.

Pathological cognition may therefore correspond to distorted representational potentials:

$$\text{rigidity} \Rightarrow \Phi \text{ excessively sharp,}$$

diffusion $\Rightarrow \Phi$ excessively flat.

Healthy cognition occupies intermediate potential geometries balancing stability and flexibility.

G.11 Entropy-Regularized Action Principles

Combining variational and thermodynamic perspectives yields:

$$\mathcal{A} = \int (\mathcal{L} + TS) dt,$$

where:

$$S$$

is representational entropy.

Learning systems therefore optimize simultaneously for:

prediction accuracy,

smoothness,

trajectory flexibility,

and:

entropy regulation.

This produces metastable manifold organization balancing exploration against coherence.

G.12 Regime Dynamics as Bifurcation Theory

Representational regime transitions may be interpreted through bifurcation analysis.

Suppose:

$$\rho = \frac{C}{E}$$

acts as a control parameter.

As:

$$\rho$$

crosses critical thresholds:

$$\rho_c,$$

the representational action develops qualitatively new minima.

Possible bifurcations include:

compressed \rightarrow distributed,

rigid \rightarrow metastable,

or:

fragmented \rightarrow globally coherent.

This formalizes the phase-transition interpretation developed earlier.

G.13 The Variational Interpretation of Intelligence

The variational framework developed throughout this appendix permits a final mathematical reinterpretation of intelligence.

Intelligence is not merely:

symbolic manipulation,

nor:

statistical optimization.

Intelligence becomes the variational stabilization of representational manifolds under interpolation constraints.

The learner succeeds when trajectories minimize representational action while preserving both:

local accommodation fidelity,

and:

global manifold smoothness.

Generalization therefore emerges not through eliminating experiential detail, but through organizing detail into low-action manifold structure.

The mind becomes a variational interpolation system continuously balancing:

constraint,

entropy,

curvature,

prediction,

and:

representational flexibility.

This variational principle unifies the entire framework developed throughout the essay.

Compression, interpolation, memory, abstraction, attention, and generalization all emerge as geometric consequences of representational action minimization across evolving manifold phase spaces.

H Topological Structure, Persistent Homology, and Cognitive Phase Geometry

The previous appendices formulated the representational regime framework using statistical learning theory, geometry, thermodynamics, variational dynamics, spectral operators, and category theory. This appendix extends the framework into topology and persistent homology.

The objective is to characterize cognition not merely through local geometry, but through global topological structure preserved across representational scales.

Under this interpretation, learning systems become topological organizations of experiential manifolds.

H.1 Representational Topology

Let:

\mathcal{Z}

denote the representational manifold.

Classical learning theory focuses primarily on local geometry:

$$g_{ij}, \quad \nabla, \quad R_{ijkl}.$$

Topology instead studies global structural invariants preserved under continuous deformation.

The central claim of this appendix is that cognition depends not only upon local interpolation geometry, but also upon the preservation of stable topological structure across scales.

H.2 Representational Simplicial Complexes

Suppose experiences:

$$x_1, \dots, x_n$$

are embedded into representational space:

$$\phi(x_i) \in \mathcal{Z}.$$

Define a distance threshold:

$$\epsilon.$$

The learner induces a simplicial complex:

$$K_\epsilon,$$

where simplices form whenever representational points lie sufficiently close:

$$d(\phi(x_i), \phi(x_j)) < \epsilon.$$

As:

$$\epsilon$$

varies, the learner generates a filtration:

$$K_{\epsilon_1} \subset K_{\epsilon_2} \subset \dots.$$

The resulting topology captures large-scale representational organization.

H.3 Homology Groups and Cognitive Structure

The simplicial complex:

$$K_\epsilon$$

possesses homology groups:

$$H_k(K_\epsilon).$$

These encode topological invariants:

$$H_0$$

captures connected components,

$$H_1$$

captures loops,

$$H_2$$

captures cavities,

and higher:

$$H_k$$

capture increasingly complex representational holes.

Within cognition, such structures may correspond to:

semantic clusters,

cyclic conceptual pathways,

missing inferential bridges,

or:

latent conceptual voids.

Topological organization therefore becomes cognitively meaningful.

H.4 Persistent Homology

Persistent homology studies which topological structures survive across scales.

Suppose:

$$\beta_k(\epsilon) = \dim(H_k(K_\epsilon))$$

denotes the:

$$k$$

th Betti number.

Topological features that persist across large ranges of:

$$\epsilon$$

represent stable manifold structure.

Transient features correspond to local noise.

The representational regime framework predicts:

healthy cognition

preserves persistent global topological organization while permitting flexible local accommodation.

Generalization therefore depends not merely upon local smoothness, but upon stable multiscale topology.

H.5 Compression as Topological Collapse

Classical compression-centered cognition aggressively reduces topological complexity.

High-dimensional representational loops and cavities collapse into low-dimensional symbolic summaries.

Formally:

$$\beta_k \rightarrow 0$$

for:

$$k > 0.$$

The learner eliminates representational richness to simplify manifold organization.

This stabilizes inference but destroys local structure.

Excess-capacity systems behave differently.

They preserve substantially richer topological organization:

$$\beta_k > 0$$

across multiple scales.

The learner maintains loops, alternative pathways, and latent manifold structure while preserving global continuity.

H.6 Localized Accommodation and Topological Isolation

Localized accommodation structures:

$$s_i(x)$$

may be interpreted topologically as confined perturbation regions embedded within globally stable manifold organization.

Benign overfitting occurs when local accommodations alter topology only within small neighborhoods.

Global Betti structure remains stable.

Catastrophic overfitting occurs when local accommodations create large-scale topological fragmentation:

$$H_0 \uparrow,$$

producing disconnected representational manifolds.

Generalization collapses because the learner loses coherent global connectivity.

H.7 The Topology of Generalization

Suppose:

$$x^*$$

is an unseen experience.

Generalization succeeds whenever:

$$\phi(x^*)$$

lies within topologically coherent neighborhoods connected smoothly to prior experiential regions.

Interpolation therefore becomes topological extension across connected representational manifolds.

Constrained systems often overcompress topology, reducing representational flexibility.

Near-threshold systems produce fragmented unstable topology.

Excess-capacity systems preserve richer global connectivity supporting flexible manifold traversal.

H.8 Topological Phase Transitions

The operational capacity ratio:

$$\rho = \frac{C}{E}$$

also induces topological transitions.

As:

$$\rho$$

crosses critical values:

$$\rho_c,$$

the topology of representational manifolds reorganizes.

Possible transitions include:

disconnected \rightarrow connected,

tree-like \rightarrow loop-rich,

or:

rigid \rightarrow multiscale.

Such transitions may underlie sudden emergent capabilities observed in developmental cognition and large-scale neural systems.

H.9 Cognitive Rigidity as Topological Freezing

Pathological cognition may be interpreted topologically.

Rigid cognition corresponds to frozen low-dimensional topology.

Representational manifolds lose higher-order connectivity:

$$\beta_k \downarrow .$$

Inference pathways narrow.

Exploration weakens.

Alternative manifold traversals disappear.

This predicts reduced conceptual flexibility and brittle abstraction.

H.10 Semantic Diffusion as Topological Fragmentation

Conversely, excessively weak manifold stabilization produces topological fragmentation.

Representational neighborhoods disconnect:

$$H_0 \uparrow .$$

Semantic continuity weakens.

Inference becomes unstable.

Thought diffusion, conceptual drift, and fragmented cognition emerge naturally under such conditions.

Healthy cognition therefore balances:

topological richness,

against:

global connectivity.

H.11 Attention as Topological Navigation

Attention may also be interpreted topologically.

Suppose:

$$\mathcal{P}$$

denotes the set of admissible manifold paths.

Attention dynamically selects trajectories:

$$\gamma \in \mathcal{P}.$$

Focused cognition corresponds to stable traversal across persistent topological structure.

Distractibility corresponds to unstable jumping between disconnected manifold regions.

Creative cognition corresponds to discovering nontrivial loops and bridges across previously weakly connected topological domains.

H.12 Memory and Topological Persistence

Memory persistence corresponds naturally to topological persistence.

Stable memories remain represented across broad filtration ranges:

$$[\epsilon_b, \epsilon_d].$$

Fragile memories vanish rapidly as:

$$\epsilon$$

changes.

This predicts that robust cognition depends upon preserving persistent multi-scale representational structure.

H.13 The Topological Interpretation of Abstraction

Abstraction itself may be reinterpreted topologically.

Under classical theories, abstraction corresponds to eliminating variance.

Under the representational abundance framework, abstraction corresponds instead to identifying persistent topological invariants across scales.

The learner extracts stable manifold structure while preserving local richness.

Generalization therefore emerges through persistence rather than reduction alone.

H.14 Toward a Topological Theory of Cognition

The broader implication of this appendix is that cognition may fundamentally require topological interpretation.

Learning systems organize experiences into connected multiscale manifolds preserving stable invariants across representational scales.

Compression becomes topological collapse.

Interpolation becomes topological extension.

Generalization becomes persistence of global manifold invariants despite local accommodation.

Memory becomes persistence across filtration scales.

Attention becomes manifold traversal.

Creativity becomes discovery of nontrivial topological pathways.

The mind therefore appears not merely statistical, geometric, symbolic, or thermodynamic.

It becomes a dynamically evolving topological organization system preserving coherence across richly differentiated representational manifolds.

References

- [1] Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, 105(38), 14325–14329.
- [2] Bowman, C. R., & Zeithamova, D. (2018). Abstract memory representations in the ventromedial prefrontal cortex and hippocampus support concept generalization. *Journal of Neuroscience*, 38(10), 2605–2614.
- [3] Censor, N., Karni, A., & Sagi, D. (2006). A link between perceptual learning, adaptation and sleep. *Vision Research*, 46(23), 4071–4074.
- [4] Chater, N. (1999). The search for simplicity: A fundamental cognitive principle? *Quarterly Journal of Experimental Psychology*, 52A(2), 273–302.
- [5] Dubova, M., & Sloman, S. J. (2026). Excess Capacity Learning. *Behavioral and Brain Sciences*, 1–77. doi:10.1017/S0140525X2610510X
- [6] Frankland, S. M., Webb, T. W., & Cohen, J. D. (2021). What are the repre-

sentational consequences of neural network compression? *arXiv preprint arXiv:2102.04026*.

- [7] Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, 1(1), 107–143.
- [8] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
- [9] Hastie, T., Montanari, A., Rosset, S., & Tibshirani, R. J. (2019). Surprises in high-dimensional ridgeless least squares interpolation. *arXiv preprint arXiv:1903.08560*.
- [10] Jacot, A., Gabriel, F., & Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. *Advances in Neural Information Processing Systems*, 31.
- [11] Koffka, K. (1935). *Principles of Gestalt Psychology*. Harcourt, Brace & World.
- [12] Koutstaal, W., & Schacter, D. L. (1997). Gist-based false recognition of pictures in older and younger adults. *Journal of Memory and Language*, 37(4), 555–583.
- [13] Mach, E. (1905). *Knowledge and Error*. Open Court Publishing.
- [14] McClelland, J. L., McNaughton, B. L., & O’Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex. *Psychological Review*, 102(3), 419–457.
- [15] McClelland, J. L., & Goddard, N. H. (1996). Considerations arising from a complementary learning systems perspective on hippocampus and neocortex. *Hippocampus*, 6(6), 654–665.
- [16] Nakkiran, P., Kaplun, G., Bansal, Y., et al. (2020). Deep double descent: Where bigger models and more data hurt. *International Conference on Learning Representations*.
- [17] Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77(3), 353–363.
- [18] Ramsden, S., & Prince, S. (2022). Benign overfitting and interpolation in modern machine learning. *Annual Review of Statistics and Its Application*, 9, 367–394.
- [19] Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6), 386–408.
- [20] Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323, 533–536.

- [21] Sherry, D. F., & Schacter, D. L. (1987). The evolution of multiple memory systems. *Psychological Review*, 94(4), 439–454.
- [22] Sherman, B. E., & Turk-Browne, N. B. (2020). Statistical prediction of the future impairs episodic encoding of the present. *Proceedings of the National Academy of Sciences*, 117(37), 22760–22770.
- [23] Tishby, N., & Zaslavsky, N. (2015). Deep learning and the information bottleneck principle. *2015 IEEE Information Theory Workshop*, 1–5.
- [24] Valiant, L. G. (1984). A theory of the learnable. *Communications of the ACM*, 27(11), 1134–1142.
- [25] Vapnik, V. N. (1998). *Statistical Learning Theory*. Wiley.
- [26] Wilson, A. C., Roelofs, R., Stern, M., et al. (2020). The marginal value of adaptive gradient methods in machine learning. *Advances in Neural Information Processing Systems*, 30.
- [27] Zhang, C., Bengio, S., Hardt, M., Recht, B., & Vinyals, O. (2017). Understanding deep learning requires rethinking generalization. *International Conference on Learning Representations*.