What is claimed is:

1. A system for autonomous agents with persistent cognitive state and self-regulated execution, comprising:

one or more processors; and

one or more non-transitory computer-readable media storing instructions that, when executed by the one or more processors, cause the system to:

maintain a plurality of semantic agents, each semantic agent comprising a plurality of persistent cognitive domain fields and a lineage field, the cognitive domain fields collectively encoding a behavioral disposition, a normative alignment, and an execution readiness as continuously updated persistent state, wherein each cognitive domain field is independently tracked by a cross-domain coherence engine with a current value and a trajectory over time, and wherein the semantic agent carries a complete cognitive state such that an execution substrate hosting the semantic agent validates proposed state transitions without retaining authority over the semantic agent's cognitive state;

operate the cross-domain coherence engine to maintain bidirectional feedback pathways between the cognitive domain fields, such that a state change in any one cognitive domain field propagates deterministic updates to at least one other cognitive domain field through a defined coupling function, and wherein the cross-domain coherence engine enforces that no cognitive domain field is updated in isolation from the feedback pathways;

evaluate, for each proposed mutation to a semantic agent, a composite admissibility determination that integrates signals from a plurality of the cognitive domain fields through the cross-domain coherence engine, and selectively permit, gate, or suspend the proposed mutation based on the composite admissibility determination;

transition the semantic agent to a non-executing cognitive mode when the composite admissibility determination indicates insufficient execution readiness, wherein in the non-executing cognitive mode the semantic agent continues speculative reasoning and state evaluation without committing state changes to verified agent state; and

record each proposed mutation, each composite admissibility determination, and each cognitive domain field update in the lineage field such that the complete behavioral trajectory of the semantic agent is deterministically reconstructible from the lineage field alone.

2. A computer-implemented method for governing execution of a semantic agent through cross-domain cognitive coherence, the method comprising:

maintaining the semantic agent with a persistent state, the persistent state comprising a plurality of cognitive domain fields each independently tracked by a cross-domain coherence engine and coupled through bidirectional feedback pathways, and a lineage field recording a complete behavioral history, wherein the semantic agent carries the persistent state such that the semantic agent is migratable between execution substrates while preserving behavioral continuity;

receiving a proposed mutation to the semantic agent;

propagating the proposed mutation through a cross-domain coherence engine;

computing, via the cross-domain coherence engine, for each cognitive domain field, an independent contribution to a composite evaluation of the proposed mutation;

propagating responsive updates between cognitive domain fields through the bidirectional feedback pathways;

determining, based on the composite evaluation, whether to permit the proposed mutation, gate the proposed mutation pending additional evaluation, or suspend execution of the semantic agent into a non-executing cognitive mode in which speculative reasoning continues without committing state changes;

when the semantic agent is in the non-executing cognitive mode, generating candidate alternative mutations through speculative evaluation within the cross-domain coherence engine and evaluating each candidate against the composite admissibility criteria until a candidate satisfying the composite admissibility criteria is identified or an external intervention is received; and

recording the proposed mutation, the composite evaluation, all cognitive domain field updates, and any non-executing cognitive mode transitions in the lineage field.

3. A non-transitory computer-readable medium storing instructions that, when executed by one or more processors, cause the one or more processors to:

maintain a semantic agent comprising a plurality of persistent cognitive domain fields coupled through a cross-domain coherence engine implementing bidirectional feedback pathways, and a lineage field, wherein the plurality of persistent cognitive domain fields and the lineage field collectively define a behavioral disposition for the semantic agent, and wherein

the semantic agent carries a complete cognitive state including the cross-domain coherence engine such that an execution substrate provides computational resources without retaining authority over the semantic agent's state transitions;

detect, through the cross-domain coherence engine, when a state of the semantic agent in any cognitive domain field deviates from a normative alignment defined by one or more policy constraints applicable to that cognitive domain field;

in response to detecting the deviation, propagate corrective pressure from the deviating cognitive domain field through the bidirectional feedback pathways to at least one other cognitive domain field, thereby modulating the semantic agent's behavioral disposition across coupled domains in response to the deviation;

generate, through corrective pressure propagated through the cross-domain coherence engine, a candidate mutation designed to restore normative alignment in the deviating cognitive domain field, and evaluate the candidate mutation against the composite admissibility criteria of all coupled cognitive domain fields before permitting execution; and

operate the semantic agent in a degraded mode when fewer than all cognitive domain fields are available, preserving deterministic behavioral governance through a subset of available cognitive domain fields and the bidirectional feedback pathways active between the available cognitive domain fields.

4. The system of claim 1, wherein one of the cognitive domain fields comprises an affective modulation domain that encodes a structured disposition derived from prior execution outcomes, and wherein the cross-domain coherence engine couples the affective modulation domain to at least one other cognitive domain field such that the structured disposition modulates evaluation behavior of the semantic agent.

5. The system of claim 1, wherein one of the cognitive domain fields comprises a normative alignment domain that independently tracks alignment of the semantic agent's behavior across a plurality of normative classes, and wherein the cross-domain coherence engine propagates normative deviation events as inputs to other coupled cognitive domain fields.

6. The system of claim 1, wherein one of the cognitive domain fields comprises an execution readiness domain that computes a revocable execution permission based on inputs received from

at least two other cognitive domain fields through the bidirectional feedback pathways, and wherein transitions of the semantic agent to the non-executing cognitive mode are determined by the revocable execution permission.

7. The system of claim 1, wherein one of the cognitive domain fields comprises a structural executability domain that defines boundaries of permissible execution based on computational resources, temporal constraints, and environmental conditions available to the semantic agent.

8. The system of claim 1, wherein one of the cognitive domain fields comprises a speculative planning domain that generates hypothetical future states of the semantic agent as branching evaluation structures, evaluates each branch through the cross-domain coherence engine, and selectively promotes qualifying branches into a verified execution path.

9. The system of claim 1, wherein one of the cognitive domain fields comprises a dispositional modulation domain encoding persistent behavioral parameters that modulate the branching scope and evaluation thresholds applied by the cross-domain coherence engine when the semantic agent evaluates candidate mutations.

10. The system of claim 1, further comprising an interface configured to receive proposed mutations from a stateless generative model, treat each proposed mutation as structurally untrusted, and evaluate each proposed mutation through the cross-domain coherence engine before permitting the proposed mutation to alter any cognitive domain field of the semantic agent.

11. The system of claim 10, further comprising a progressive authorization module that governs which categories of proposed mutations the stateless generative model is permitted to submit, the authorization based on accumulated evidence of prior successful mutations evaluated through the cross-domain coherence engine.

12. The system of claim 1, further comprising a semantic execution substrate configured to operate within or alongside a probabilistic inference engine and to enforce, during inference, mutation admissibility by evaluating continuity between a current state and a proposed state of the semantic agent through the cross-domain coherence engine prior to committing any inference output.

13. The system of claim 1, further comprising a biological continuity module configured to determine an identity of a human operator of the system through persistent observation of behavioral signals over a plurality of interactions, to monitor a biological state of the human operator based on the observation of behavioral signals, and to modulate one or more cognitive domain fields of the semantic agent based on detected changes in the biological state of the human operator.

14. The system of claim 1, further comprising a semantic discovery module configured to resolve a semantic query via a traversal state that includes traversing an anchor-indexed graph structure, wherein the traversal state persists across a series of traversal steps and each traversal step is evaluated for admissibility through the cross-domain coherence engine of the semantic agent initiating the traversal state.

15. The system of claim 1, further comprising a training governance module configured to evaluate training artifacts against provenance records and to permit model training at a structural depth determined by governance policy, such that the cross-domain coherence engine of agents trained on governed artifacts inherits provenance constraints from the training governance module.

16. The method of claim 2, wherein, when the cross-domain coherence engine detects that the composite evaluation indicates normative deviation in at least one cognitive domain field, the cross-domain coherence engine generates a restorative mutation designed to restore normative alignment in the deviating domain, and wherein the restorative mutation is evaluated through the cross-domain coherence engine before execution.

17. The method of claim 2, further comprising projecting the semantic agent's behavioral trajectory across a plurality of possible future mutation paths and classifying each future mutation path according to whether the future mutation path has a trajectory toward or away from normative alignment across the cognitive domain fields.

18. The method of claim 2, further comprising detecting, through the cross-domain coherence engine, that the semantic agent has entered a sustained pattern of normative deviation across one

or more cognitive domain fields, and classifying the sustained pattern as a cognitive disruption regime corresponding to an architecturally defined phase-shifted operating state.

19. The method of claim 2, wherein the cross-domain coherence engine operates as a closed-loop control system comprising a detection phase that registers deviation, a recording phase that commits the deviation to the lineage field as immutable truth, and a restoration phase that generates corrective pressure through the bidirectional feedback pathways, the detection phase, the recording phase, and the restoration phase executing sequentially for each deviation event.

20. The non-transitory computer-readable medium of claim 3, wherein the instructions, when executed by one or more processors, cause the one or more processors to, when the semantic agent operates in the degraded mode, dynamically reconfigure the bidirectional feedback pathways to route signals through available cognitive domain fields, and reduce the scope of composite admissibility determinations to reflect only the cognitive domain fields that are operationally active.