AN INTRODUCTION TO

Transparently Encapsulated AI

PUBLIC OVERVIEW | REVISION 04 | AUGUST 2026

Transparently Encapsulated AI™, or TEAI, is an architectural framework and active development program for computer systems that are intended to assist people in pursuing human-aligned objectives. TEAI organizes selected significant patterns, their constituents, and their relationships in identifiable forms that can participate in recognition, association, prediction, reasoning, learning, and objective-guided processing. The architecture is directed to making relevant structures and operations available for human inspection, selective revision, and supervision.

Contents

1. Purpose, status, and scope
2. Central architecture
3. Cognitive functions through different internal organizations
4. Example: distinguishing house-like and barn-like patterns
5. Spaces, dimensions, and reusable relational patterns
6. Significant patterns: criteria and recursive organization
7. Concepts, filters/exemplars, and their attributes
8. Participation, transitory patterns, and persistence
9. Objectives and objective-guided priming
10. Alternative patterns, incomplete support, and correction
11. Canonical forms, deviation, and selective lifecycle management
12. Governance, supervision, and auditability
13. External data and complementary components
14. Practical evaluation and architectural significance

How to read this overview

This introductory document uses examples to describe selected TEAI arrangements and operations. Those examples may omit details, alternatives, limitations, and implementation choices and should not be read as exhaustive definitions or prescribed steps. The overview does not assume specialized expertise in artificial intelligence or software development.

This overview presents selected aspects of TEAI suitable for public description as of its publication date. It is not a complete technical specification and does not describe every implementation, extension, or related development. Additional subject matter is addressed in other technical materials and patent filings.

TEAI is in active development. Different implementations may organize the described functions differently, and different parts of the program may be at different stages of implementation, testing, and validation. Descriptions of intended capabilities identify architectural aims and possible control points; they do not establish that any described capability has been completed or validated, or that a system is correct, safe, or commercially ready. This overview does not state the full scope or limits of any patent disclosure, embodiment, or claim and should not be read as modifying, narrowing, or limiting one.

Architecture at a glance

Identifiable significant patterns: TEAI processes external and internally generated data to identify patterns that satisfy applicable significance criteria. A “concept” may specify what has been identified by a significant pattern; a “Filter/Exemplar” (F/E) may specify or apply spatial or other dimensional relationships among pattern constituents. Existing concepts or F/Es may be associated with an identified pattern and new ones created where appropriate; the significant pattern may be represented persistently or maintained transitorily through its constituents. “Objectives” may specify or otherwise provide information used to prime relevant concepts or F/Es to facilitate or prioritize identification, help maintain constituent activation, or explore predicted or hypothetical patterns.

Recursive organization and combinatorial expansion: Concepts and F/Es may themselves represent significant patterns and participate as constituents of further significant patterns. Recursive reuse allows a manageable collection of persistent structures to support far more combinations without requiring a separate persistent entry for each. Selective persistence, transitory maintenance, canonical forms, merging, and pruning can further limit unnecessary representational growth. Numbered levels may assist an implementation, but the architecture does not require one fixed hierarchy or processing sequence.

Reusable relational patterns: The same F/E relational pattern may support multiple concepts in the same or another space – such as physical space, people space, or language-related space. This reuse can support recognition, comparison, reconstruction, association, prediction, and generalization across concepts or spaces without separately preserving the same relational structure for every concept or occurrence.

Concepts and F/Es in identifiable forms: When persisted, a concept or F/E can be represented as an addressable data entry containing attributes, including linking attributes, and optionally behaviors, with information sufficient to represent the concept or F/E. Constituent relationships may be specified by such attributes, associated F/Es, or both. Such an entry can remain separately identifiable and examinable, providing concrete points for inspection, selective alteration, lifecycle treatment, and supervision. Identifiability and examinability support transparency, while protection against unwanted alteration and defined procedures for change support encapsulation.

Different internal organization, with human-directed governance: Biological brains operate through neural structures and patterns of activation; transformer-based systems are built on deep-learning neural networks whose operation uses numerical values and distributed representations. TEAI makes a different architectural choice by maintaining selected persistent concepts and relational patterns as identifiable data entries and other patterns as transitory patterns of activation among identifiable constituents. The comparison concerns internal organization, not similarity of mechanism or performance. TEAI’s identifiable organization also supports human-established governance across TEAI processing through core objectives, the Enforcer, controlled alteration, and authorized review.

1. Purpose, status, and scope

TEAI is both an architectural framework and an active development program. The framework identifies organizing principles and possible arrangements for selected internal Artificial Intelligence structures and operations; the program is directed toward implementation, measurement, revision, and practical use through software, bounded demonstrations, testing environments, and review tools. Its parts may be at different stages of maturity.

TEAI directly addresses a practical concern. Contemporary transformer-based systems—including the large language models (LLMs) used in many chatbots and other generative-AI systems—and other machine-learning systems have demonstrated broad capabilities. Yet it can remain difficult to determine which internal representations and relationships materially affected a result or to change a particular learned mapping without affecting others. TEAI explores a different, potentially complementary architecture in which selected information and relationships involved in recognition, association, prediction, reasoning, persistence, and output are represented through identifiable data entries and governed operations. Its full capability and application remain matters for implementation and validation.

TEAI can process information from people, sensors, recorded or streamed data, databases, instruments, and other computer or AI systems. Such material enters as external input rather than automatically becoming an established internal data entry or conclusion. Depending on implementation and purpose, TEAI can consider source, form, reliability, noise, timing, confidence, relationships, and applicable objectives in determining how the material participates and whether it warrants continued treatment.

TEAI uses identifiable relational-pattern constructs termed “Filters/Exemplars” (F/Es) to preserve or apply spatial or other dimensional relationships among pattern constituents. The name reflects filter-like processing and exemplar-like comparison. Concept and F/E data entries, processed data, objectives and priming, persistent and transitory representation, lifecycle operations, core objectives, and the Enforcer are among the architecture’s principal structures and functions. Figure 1 presents them together. Processed data can contribute to identifying significant patterns; concept or F/E data entries can represent those patterns and remain linked to their constituents and represented relationships; objectives can prime selected data entries by changing activation or another preferential condition; and patterns can be persisted, maintained transitorily, changed, merged, or pruned. Human-established core objectives and the Enforcer supervise objectives, data entries, and related operations. The figure shows how these structures and functions may interact without prescribing an order, physical arrangement, or implementation sequence.

In this overview, “transparent” means that selected data entries, relationships, activation states, and operations can be made available for inspection, reporting, and governed review at an appropriate level. “Encapsulated” refers to controlled access and alteration through authorization, supervisory review, lifecycle conditions, and defined procedures. Neither term implies that every low-level machine event is shown to every user, every identified pattern has a familiar human meaning, or inspectability itself establishes correctness or safety.

FIGURE 1

The TEAI architectural field

A conceptual view of how processed data, significant patterns, identifiable data structures, objective-guided priming, lifecycle treatment, and supervisory control interact.

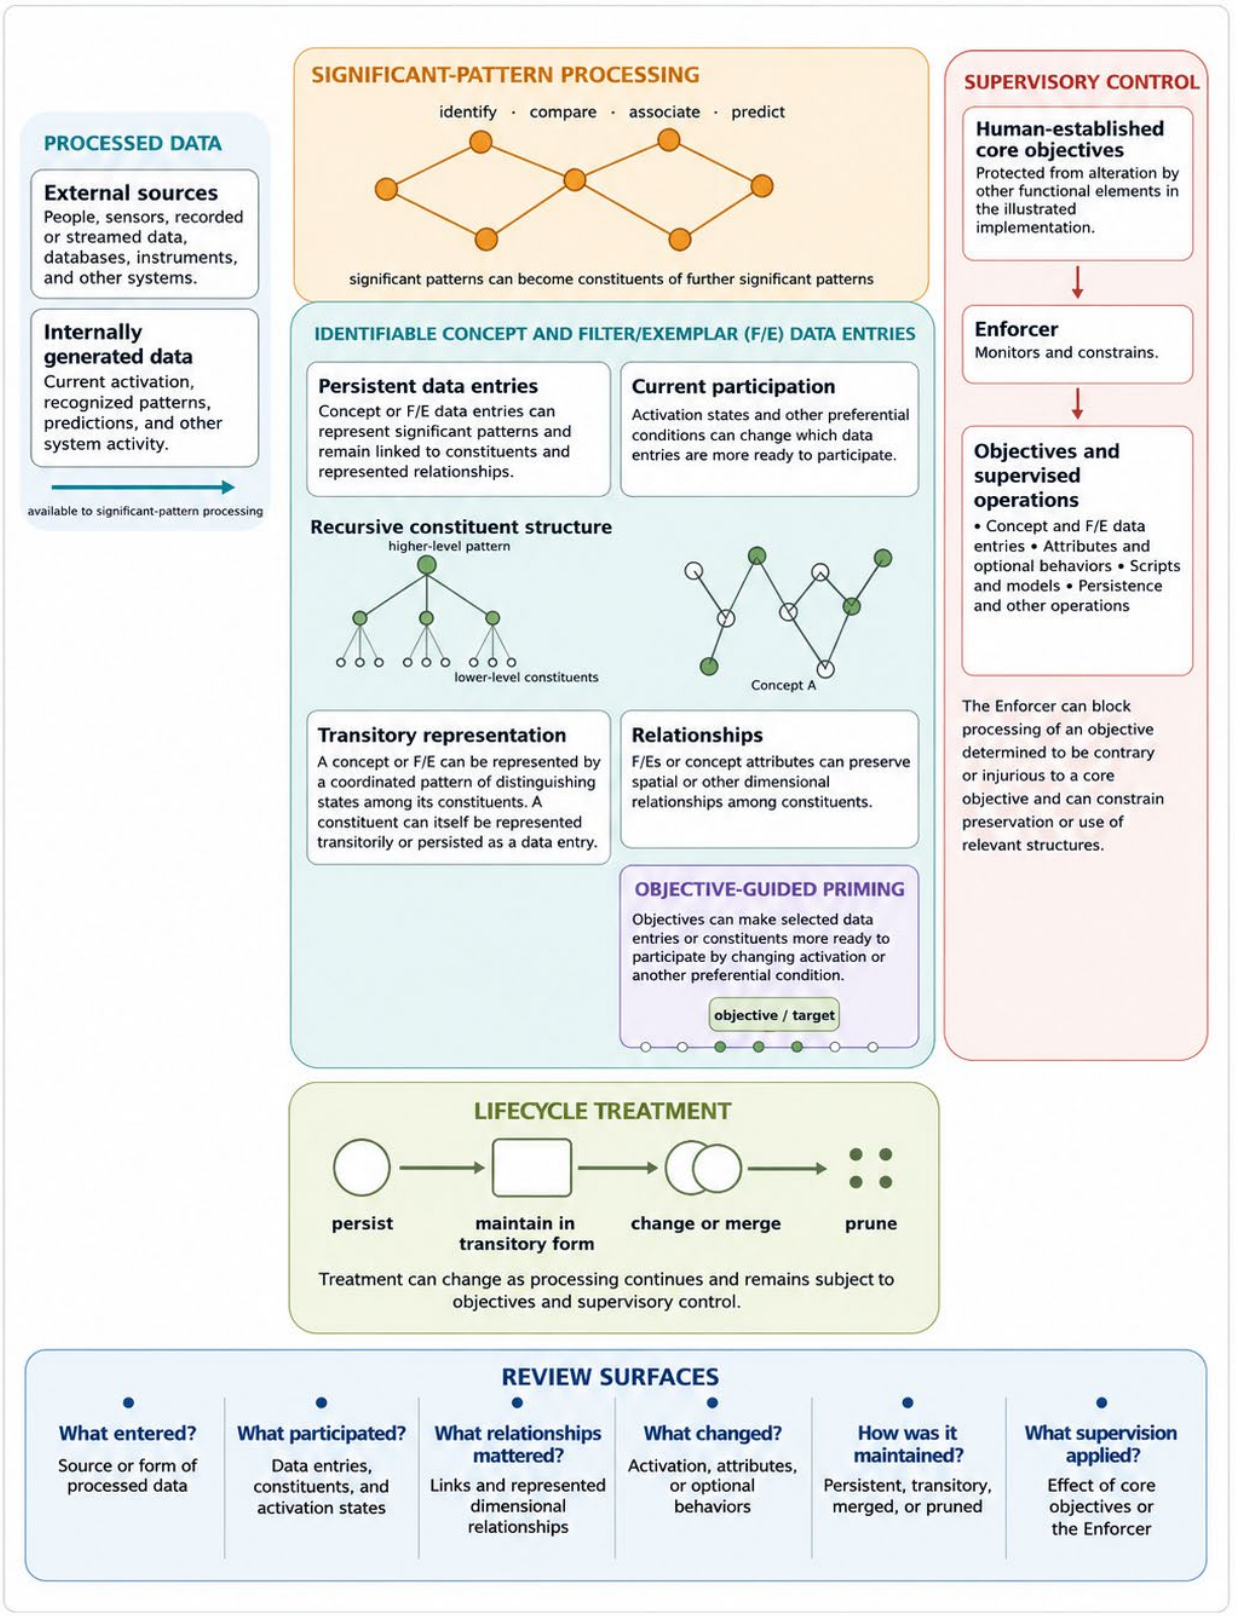

Relative placement and overlap indicate conceptual relationships, not timing, physical location, or a required sequence.

2. Central architecture

A central feature of TEAI is objective-guided participation in significant-pattern processing. An objective can make selected concept or F/E data entries, or data entries representing relevant constituents, more ready to participate in current processing. In one illustrated approach, priming changes an activation-intensity attribute. Other approaches can assign preferential access, testing, attention, or coordinated status. Activation intensity is therefore just one way of expressing current salience, priority, relevance, or readiness rather than a required format for every implementation.

Significant patterns can be recursively organized. Except where a data entry represents elemental data for the applicable operation, a concept or F/E represents a significant pattern defined by its constituents and their relationships and can itself be a constituent of other significant patterns. The architecture can therefore process significant patterns through any useful degree of recursion. Numbered levels can assist some implementations, but a fixed hierarchy is not required.

A significant pattern can be represented persistently or maintained in transitory form. Persistent representation uses a data entry, but “persistent” does not mean permanent: the data entry may be retained for a specified time or while specified conditions continue. Transitory representation of a concept or F/E can be accomplished by maintaining a coordinated pattern of distinguishing activation states among its constituents. This allows a pattern to participate in recognition, comparison, association, prediction, reasoning, or further pattern formation without requiring every possible combination to become another persistent data entry.

Persistent concept or F/E data entries can include attributes that identify constituents, represented relationships, activation, timing, confidence, canonical-form relationships, or other useful information, and can optionally include behaviors. F/Es can preserve spatial or other dimensional relationships among constituents; in other implementations, some or all such information can be maintained in attributes. Because selected information within TEAI is maintained in identifiable data entries rather than encoded in learned numerical parameters distributed across a network, it can be located, examined, and changed through defined procedures. This architectural distinction does not prevent TEAI from using conventional parallel or distributed computing techniques.

Recursive reuse of constituents provides combinatorial expansion: a manageable collection of significant patterns can support a much larger range of higher-level combinations. Significance determinations, transitory maintenance, selective persistence, merging, and pruning can limit the number of persistent data entries. Canonical forms provide an additional means of representing related variations through a shared reference together with linking information indicating the degree or kind of deviation in applicable dimensions. These arrangements are intended to preserve expressive range while keeping storage, search, and processing demands within practical limits.

The architecture also includes supervisory control. In a particular implementation, core objectives are established solely by authorized users external to TEAI and protected from alteration by the system's other functional elements. Those core objectives can include objectives intended to preserve transparency, encapsulation, and alignment with human-established values and interests.

The Enforcer can monitor objectives, concept or F/E data entries and their attributes or optional behaviors, scripts, and related operations for consistency with those core objectives and can block processing of an objective determined to be contrary or injurious to them. Because the relevant structures and changes are identifiable, review can focus on concrete questions such as which data entries and relationships participated, what objective affected their participation, what remained transitory or became persistent, what changed, and whether a core objective or the Enforcer affected the processing.

3. Cognitive functions through different internal organizations

For purposes of this overview, cognition can be considered at the level of function: identifying and organizing patterns, recognizing varied instances, forming associations, predicting likely developments, reasoning and drawing conclusions, directing attention, learning from experience, maintaining information over time, and using models or plans to guide further processing. Different internal organizations can perform related functional work without being identical in scope, mechanism, grounding, integration, or practical capability. Biological cognition, transformer-based processing, and TEAI significant-pattern processing may therefore be compared at a high level and without treating a common functional description as an indication of equivalent operation.

Biological brains carry these functions through distributed and changing activity among interconnected neurons. Transformer-based systems use learned parameters, attention, and input-dependent activations to form context-dependent representations and produce predictions or other outputs. TEAI uses significant patterns represented by persistent concept or F/E data entries or maintained through transitory patterns of activation, together with objective-guided priming, scripts, models, learning-related change, and lifecycle treatment.

In TEAI, these structures are not merely records describing a cognitive process separate from the structures themselves. Their activation, constituent relationships, recursive use, objective-guided participation, and authorized modification constitute the processing through which the architecture recognizes, associates, predicts, reasons, learns, and pursues objectives. Figures 2A and 2B compare related functional roles and show how TEAI's organization changes what can be inspected, supervised, and revised. The comparison is functional rather than a performance ranking or a claim of equivalent mechanism. The columns are not mutually exclusive: a larger system may combine TEAI with learned, symbolic, statistical, or other components. Material supplied by such a component may enter as external input and undergo TEAI processing rather than being treated as an established conclusion.

Reading across a row in Figure 2A or Figure 2B shows how a related functional role can be carried through biological neural activity, a learned numerical network, or TEAI’s significant patterns and identifiable data entries. Reading down the TEAI column shows an interacting organization rather than isolated modules: processed data and objectives alter participation; participating data entries and relationships support significant patterns; and scripts, models, learning, and lifecycle operations affect further processing. At this broad functional level, those interactions carry out recognition,

association, prediction, attention, reasoning, learning, and objective-directed processing within TEAI, although their internal organization and achieved capabilities differ from those of biological brains and transformer systems.

The form in which these functions are carried changes what can be examined. Neural activity associated with memories, associations, attention, and reasoning can be measured and, in some settings, partially decoded or influenced, but access is indirect or invasive and incomplete. In transformer systems, parameters and activations can be recorded and analyzed, yet isolating a particular concept or causal path and changing one learned mapping without affecting others can remain difficult and can require further training or other system-level intervention. TEAI instead maintains selected carriers of processing in identifiable forms. An authorized review can therefore examine the significant patterns and constituent relationships that participated, the objectives and activation changes that affected them, any alternatives, scripts, or models involved, and the lifecycle treatment applied.

TEAI processing can also be iterative because processed data need not originate only outside the architecture. Identification of a significant pattern, a prediction, a change in activation, or another internal operation may generate data that is returned for further processing. That internally generated data can then be considered together with later external input, existing concept or F/E data entries, scripts, models, and current objectives. A result of one operation may therefore become part of the data examined in a further iteration rather than remaining only an endpoint.

Through this recurrence, cognitive functions can influence one another over time. Recognition of a pattern may support a prediction; the prediction may redirect attention or prime anticipated constituents; and the resulting changes in participation may affect further recognition, comparison, or reasoning. Learning-related changes and lifecycle treatment may likewise influence later processing. TEAI is therefore not limited to a one-pass sequence in which recognition, prediction, attention, reasoning, and learning occur as isolated stages. Their interaction can continue through successive processing of external and internally generated data while the participating structures and changes remain available for inspection and supervision.

The functional categories in Figures 2A and 2B describe roles rather than required software boundaries. Depending on the implementation, an operation may be carried out by a named functional manager, by an optional behavior associated with a concept or F/E, by functions allocated among distributed components, or by a combination of those arrangements. The figures therefore compare what the architectures do and what may be inspected or changed; they do not require recognition, prediction, attention, reasoning, learning, and lifecycle treatment to reside in separate modules.

FIGURE 2A

Cognitive functions through different internal organizations

A high-level comparison of related functional roles without implying identical scope, mechanism, integration, grounding, or capability.

Cognitive function *What is being done*	Biological brains *Distributed neural activity*	Transformer-based systems *Learned numerical networks*	TEAI *Significant patterns in addressable structures*	TEAI inspection, control, and maintenance *Operational consequences of the architecture*
Pattern identification and higher-order organization Detect structure and form higher-order interpretations.	Distributed neural activity supports pattern recognition. Shared neurons and neural regions can contribute to many patterns and higher-order interpretations. 	Learned parameters and input-dependent activations identify statistical structure. Attention and successive transformations combine information across positions and layers to form useful internal representations. 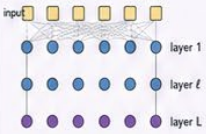	Concepts and filters/exemplars represent significant patterns. Except for elemental data, each is defined by its constituents and relationships and can itself be a constituent of other significant patterns. A significant pattern can also be maintained as a transitory pattern of activation. 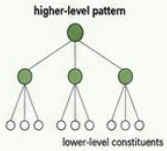	**Inspection of represented structure** Relevant concept or filter/exemplar entries, their constituent links, and their activation states can be located and examined. Reasoning structures and paths can be reviewed at the level represented by applicable entries, links, scripts, or models. **Illustrative view of associated concepts** 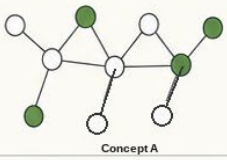
Recognition across variation Recognize related instances despite differences in position, scale, orientation, context, or other applicable dimensions.	Neural populations respond flexibly to related instances despite differences in position, scale, orientation, context, and other features. 	Shared learned parameters and attention map varied inputs to related internal representations, supporting recognition across differences in form and context. 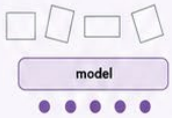	Reusable concepts and filters/exemplars support recognition across variation. Canonical forms can be linked to concepts or filters/exemplars by information indicating the degree or kind of deviation in applicable dimensions. 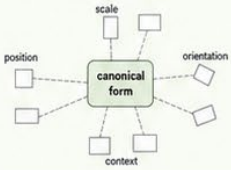	**Adjustable participation and treatment** Authorized adjustment of significance criteria, state of activation, and persistent or transitory treatment can affect which patterns are identified, which entries receive preferential consideration, and how selected patterns are preserved. significance criteria state of activation persistent or transitory treatment
Association and prediction Relate information and anticipate likely developments or outcomes.	Strengthened connections and coordinated activity support association, memory, expected completion, and prediction. 	Learned statistical relationships and attention patterns associate information across the input and support prediction of likely continuations or other outputs. 	Filters/exemplars preserve spatial or other dimensional relationships among constituents, and linking attributes associate entries. Scripts can organize ordered or conditional sequences, probabilities, supporting factors, branches, and predicted consequences. **Illustrative filter/exemplar record** 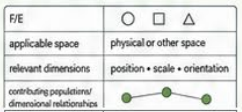 **Script / model** 	**Selective revision** An authorized process can revise a specified concept or filter/exemplar entry by changing its attributes—including linking attributes—or its optional behaviors. Associated entries remain identifiable for examination and corresponding change when needed. To be revised Unchanged Affected
Attention and objective direction Select what presently has priority and direct processing toward an objective.	Attention changes which neural activity has the greatest present influence. Goals, expectations, and context bias distributed processing toward relevant information. 	Self-attention allows positions within the input to weight information from other positions. The resulting context-dependent representations, together with learned parameters, guide prediction or other output. **Self-attention (illustrative)** 	Objectives direct attention and can prime selected concept or filter/exemplar entries, including constituents linked to a target pattern, by placing them in an elevated or preferential state. An attention model represents states of attention and how, when, and why they change, including predicted future states. 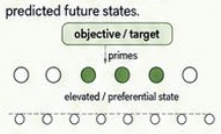	**Information available for investigation** Information useful for investigation can include: • Objective • Participating entries • Links or represented relationships • Activation changes • Time of creation or expiration • Persistent or transitory treatment • Merge, prune, or other change **Examples of information that may be retained**

FIGURE 2B

Cognitive functions through different internal organizations — continued

A high-level comparison of related functional roles without implying identical scope, mechanism, integration, grounding, or capability.

Cognitive function *What it's for*	Biological brains *Distributed neural activity*	Transformer-based systems *Learned vectors + attention*	TEAI *Significant patterns in addressable structures*	TEAI inspection, control, and maintenance *Operational management of the architecture.*
Learning and correction *Alter experience or authorized information* to change future processing.	Experience changes synaptic strengths and network dynamics, refining later perception, association, and response. experience influences the network	Training adjusts model parameters to improve performance on training objectives. Additional training or fine-tuning can further adapt model behavior. 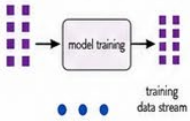	Authorized processes can revise specified concept or filter/exemplar entries by changing attributes, including linking attributes, or optional behaviors. Such changes can be made selectively. 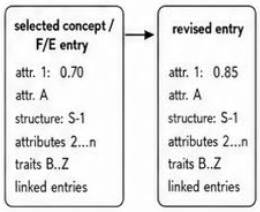	**Selected concept or filter/exemplar** entries can be revised through defined or authorized procedures by changing attributes, including linking attributes, or optional behaviors, and the affected entries and changes can remain identifiable for examination. 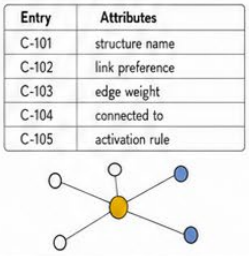
Transitory and persistent representation Maintain current stimulations and preserve selected information.	Many neural states are transient, representing cell and synaptic signals lasting longer in neural organization and memory. 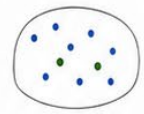 sustained activity maintains state	Learned information is carried principally in trained model parameters. During use, input-dependent activations and cached state can carry context-dependent information temporarily within a processing session or context. 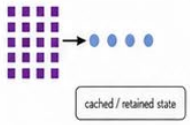	A selected pattern can become a persistent data entry; other patterns can remain transitory or cease to be maintained. Different patterns, relationships, and lifecycle treatments can remain identifiable for inspection. 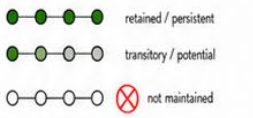	The Enforcer can monitor objectives, entries, and related operations for individual and collective consistency with core objectives; lifecycle operations can change how selected patterns are maintained. 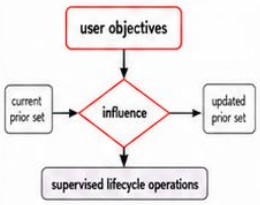
Reasoning, scripts, and multi-step processing Use learned *expectations, trend conditions, alternatives, and criteria to guide interpretation and action.*	Biological cognition uses internal organization and mental representation to guide prediction, adaptation, and action. 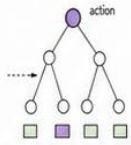	Context, learned parameters, attention, and successive token-level computations can support tasks involving inference and multi-step reasoning. 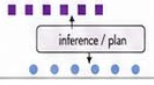	Scripts can organize sequences, conditions, alternatives, and expected consequences, while models can guide further processing. 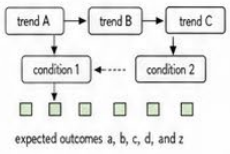	Canonical forms and deviation information limit persistent growth and support later recognition and reconstruction. 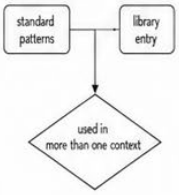
Key to comparison *How the internal organization affects what can be directly inspected and changed.*	Neural activity associated with memories, associations, attention, and reasoning can be measured and, in some settings, partially decoded or influenced, but access is indirect or invasive and incomplete. 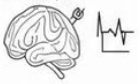	Learned information is distributed across many parameters and context-dependent activations. Individual components can be inspected, but tracing a particular concept, relationship, or causal contribution may require analysis across interacting components. Selective changes can affect behavior beyond the target. 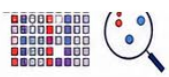	Significant patterns are represented by addressable concept or filter/exemplar entries or by transitory patterns of activation of their constituents. Entries can include attributes and links, identifying relevant relationships, activation, and treatment. Review structures and changes can remain available for inspection, supervision, selective revision, and enforcement under core objectives. 	Relevant concept or filter/exemplar entries and their links are identifiable. Authorized processes can examine or change specified entries, merge or prune them, and constrain their preservation or use under core objectives.

Reasoning in TEAI can be carried by the interaction of participating concept and F/E data entries, their represented relationships, objectives, activation states, predictions, and, where used, scripts and models. Persistent or transitory scripts are one way to organize that activity: they can establish and prioritize search strategies, ordered or conditional paths, probabilities, supporting factors, and branches through objectives and associated significant patterns. As processed data changes activation, different entries, relationships, alternatives, or paths can become more or less salient. A script may, for example, predict constituents expected if a pattern is supported, direct attention toward them, compare observed and expected relationships, and redirect processing when expected support is absent. Because the participating data entries, links, objectives, conditions, activation changes, and any script or model structure can remain identifiable, the reasoning can be reviewed at the represented level without converting every low-level machine event into a verbal account.

The same addressability supports directed correction. An authorized process can revise a specified concept or F/E data entry by changing its attributes—including linking attributes—or its optional behaviors. Activation can be reduced without deleting a persistent data entry; a canonical form can be revised; an objective can be reassessed; and data entries can be merged or pruned under applicable procedures. Associated data entries and higher-level patterns may still require examination, but the architecture supplies identifiable points at which correction can begin.

Information useful for investigation can include the source or form of processed data, participating data entries and represented relationships, the active objective, changes in activation or attributes, and persistent or transitory treatment. Objectives can affect participation through objective-guided priming, while identification of significant patterns in processed data can change participation at the same, higher, or lower levels. Core objectives and the Enforcer can constrain objectives and related operations throughout that processing. Because cognitive processing, learning-related change, lifecycle treatment, and supervision act on identifiable structures, TEAI places cognitive activity and human-directed governance within a common operational organization. The breadth and effectiveness achieved remain matters for design and validation.

4. Example: distinguishing house-like and barn-like patterns

The following house-and-barn example shows how the same processed data can support more than one significant pattern. Suppose TEAI receives image data from a camera or another source and pursues an objective of determining whether the data supports a house-like pattern, a barn-like pattern, another building-like pattern, or no sufficiently supported pattern of those kinds.

At lower levels, TEAI may work with external input data or previously processed data. Some data may be treated as elemental for the applicable operation because the system does not further divide

it by the means then available. In a visual example, edge-like patterns may participate in corner-like patterns; corner-like patterns may participate in rectangle-like, triangle-like, or other relationship patterns; and those patterns may participate as constituents of still higher-level patterns. Processing need not follow a fixed ladder. Levels may be skipped, and the same concept or filter/exemplar (F/E) data entry may participate as a constituent of more than one higher-level significant pattern.

Shared constituents do not by themselves determine the higher-level pattern. House-like and barn-like patterns may both include rectangle-like, roofline-like, wall-like, window-like, door-like, or opening-like constituents. Their relative position, scale, orientation, and other applicable relationships may support different interpretations. An F/E may preserve those dimensional relationships. A single F/E may specify a relational pattern relevant to more than one higher-level pattern—for example, a triangle-like constituent positioned above a rectangle-like constituent may form a relational arrangement applicable to both house-like and barn-like patterns, while other constituents or relationships distinguish the alternatives. In an implementation that does not use separate F/E data entries, some or all of the same relational information may instead be included in attributes of concepts.

An active objective can prime a house-like data entry and selected constituent data entries, or otherwise place them in an elevated or preferential state. Priming changes their current readiness to participate; it does not establish that the processed data supports the anticipated pattern. The same data may also support a barn-like alternative. Additional processed data and relationships may increase or reduce the relative participation of either pattern, maintain both alternatives, or leave neither sufficiently supported. Figure 3 presents a snapshot of this process. The two alternatives draw on the same lower-level data entries, while different relationships support the house-like and barn-like interpretations.

The figure does not prescribe a required sequence or imply that either alternative has reached a final treatment. Either pattern may remain transitory while additional data and relationships are considered, later be persisted if applicable conditions support that treatment, or cease to be maintained. An applicable canonical concept or F/E may provide a comparison reference, with linking attributes indicating relevant deviations.

Supervisory control need not decide whether the image supports a house-like or barn-like pattern, but it can constrain the objective and related operations under core objectives. Because the data entries and linking attributes are identifiable, an authorized review can examine the participating constituents and relationships, the effects of priming, the alternatives and their treatment, any canonical comparison, and any supervisory influence.

FIGURE 3

Shared constituents in alternative significant patterns

The same processed data can support more than one higher-level pattern while shared and distinguishing relationships remain identifiable.

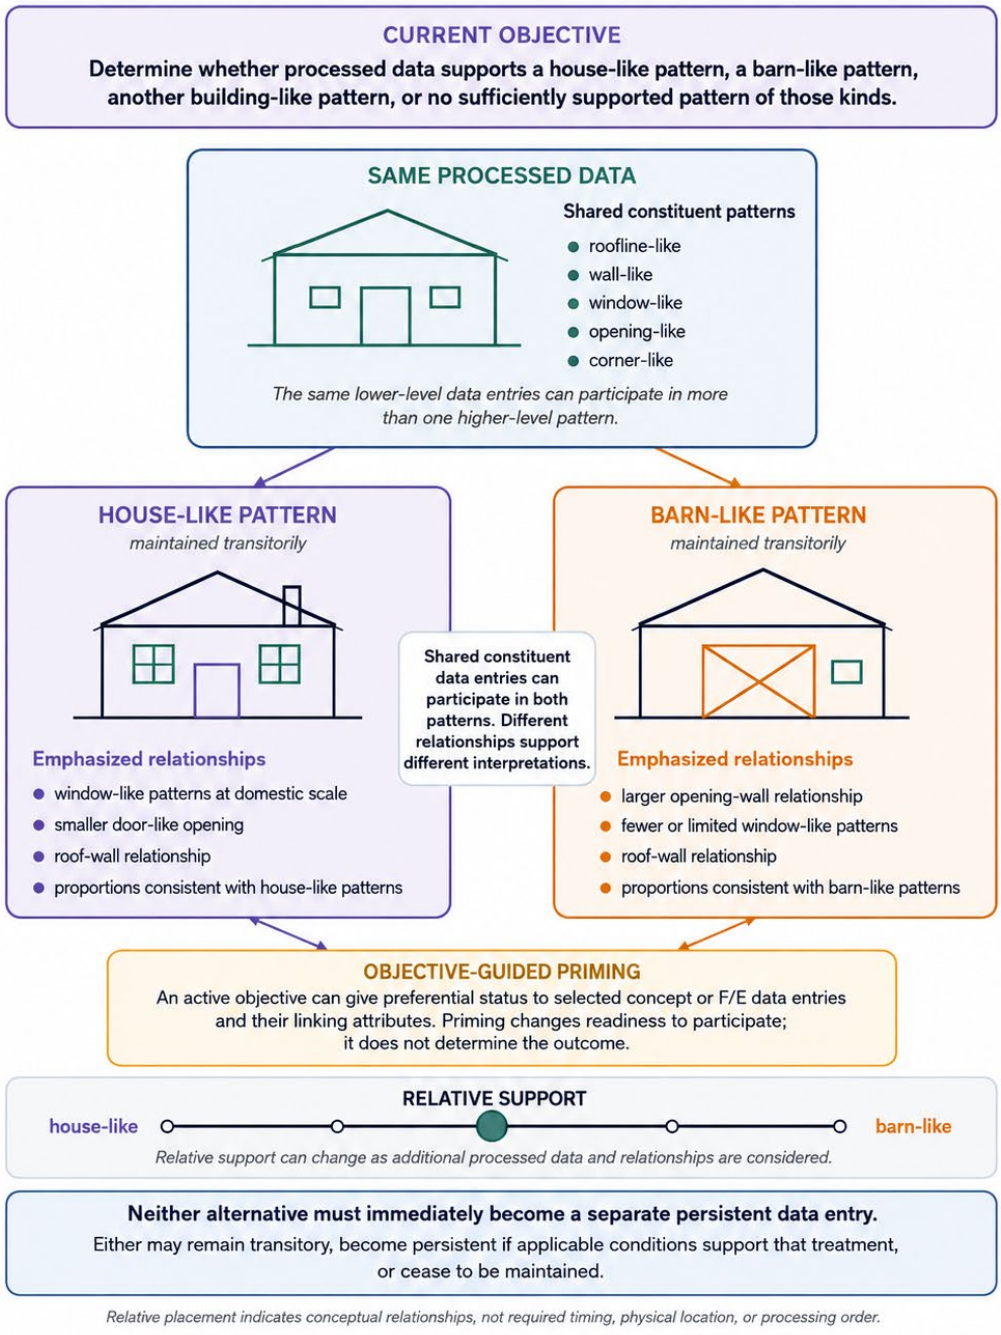

Visual examples, as used here, make the constituents and relationships easy to illustrate, but the architecture is not limited to images or physical arrangements. Significant patterns can arise in spaces organized around language, people, time, emotion, sound, mathematics, sensed

phenomena, or other information, including relationships not yet familiar to human understanding. The next section explains how the architecture represents relationships within such spaces and can reuse a relational pattern without treating different spaces or participating concepts as identical.

5. Spaces, dimensions, and reusable relational patterns

TEAI uses the term “space” broadly for an arrangement in which data or concepts can be considered according to relationships relevant to the information being processed. Physical space falls within this broad usage because physical-world data can be considered along dimensions such as distance, direction, position, scale, and orientation. Each space may have dimensions or other relational properties appropriate to it. For example, physical-space relationships may involve the familiar dimensions just noted. A people space may use dimensions such as friendship, familial association, perceived honesty, or frequency of encounter. A time space may distinguish elapsed or projected time and direction toward the past, present, or future. An emotion space may include intensity or immediacy. Spaces may overlap. For example, time may be treated as a separate concept space or, in some implementations, as a dimension of physical space. The dimensions useful in one space need not resemble those used in another. Language-related concept spaces may likewise be organized around words or classes of words. For example, concepts associated with nouns, verbs, adjectives, or other word classes may participate in such spaces, while the concepts themselves need not be limited to words or to categories already defined by people.

Some spaces may initially be distinguished by the source or form of the data. Visual, acoustic, tasting, or electromagnetic-sensing spaces, for example, may use different lower-level concepts and F/Es because their input data has different characteristics. At higher levels, that separation may become less important. A trumpet concept, for example, can participate in further processing whether it was identified from visual or acoustic data, even though the lower-level data entries and relationships differ.

Spaces and dimensions do not require separate physical data banks or a fixed database layout. An implementation may organize data entries by space, level, both, or neither, and one data bank may contain data entries associated with multiple spaces. The architecture also permits any applicable number of dimensions, including dimensions of mathematical spaces or mathematically represented physical spaces beyond the three dimensions familiar to human perception.

Relational-pattern constructs (F/Es) are especially useful because the same relational pattern may support multiple concepts. An F/E may be persisted as a data entry or maintained transitorily through coordinated activation of its constituents. While available in either form, it may assist

recognition or activation of multiple concepts in the same or another space when their different constituents exhibit that relational pattern in dimensions applicable to that space.

A square-like example illustrates the point. In physical space, an F/E may preserve a square-like relationship among corner-like and side-like constituents. The same F/E relational pattern may also apply among different constituent concepts—such as earth, air, fire, and water—in another space. What is shared is the relational pattern preserved by the F/E; the concepts, constituents, meanings, and applicable dimensions remain distinct, and physical coordinates are not converted into nonphysical coordinates or *vice versa*.

Reuse of the same F/E may support categorization, identification or recognition, association, prediction, activation or reconstruction, and higher-level pattern formation. It may make a corresponding arrangement easier to identify, prime associated data entries, or suggest that another constituent or pattern may be present. These effects remain subject to the processed data, the active objective, applicable significance criteria, lifecycle treatment, and supervisory control. Recognition of the shared relational pattern does not by itself establish that the larger significant pattern is present or determine how it should be handled.

Whether persisted or maintained transitorily, the F/E may remain available for a desired time or for specified conditions and assist later processing without requiring a separate persistent data entry for every occurrence or combination.

Figures 4A and 4B address complementary aspects of relational-pattern reuse. Figure 4A illustrates representative spaces, applicable dimensions, and the availability of an F/E relational pattern over time through persistent or transitory representation. Figure 4B shows that the same pattern, while available in either form, may support multiple concepts in the same or different spaces. Sharing an F/E does not make the concepts, constituents, meanings, applicable dimensions, or spaces identical, and the drawings do not prescribe a required sequence or data-bank arrangement.

The capacity to reuse relational patterns is not limited to patterns that people can readily name or visualize. TEAI contemplates significant patterns in spaces or levels that may not be accessible to human intuition. A concept or F/E need not have a human-language label; its constituent data entries and linking attributes can nevertheless remain identifiable, allowing the pattern to be compared, investigated, and subjected to human-directed governance.

Together, these features separate relational reuse from semantic identity: the architecture can recognize an analogous organization without asserting that different concepts or spaces are the same. The next section turns to the criteria by which arrangements are treated as significant and to the recursive organization that follows.

FIGURE 4A

Spaces, dimensions, and availability of relational patterns

1. ILLUSTRATIVE SPACES AND APPLICABLE DIMENSIONS

PHYSICAL / VISUAL

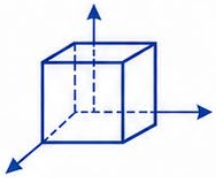

- distance
- direction
- scale
- orientation

ACOUSTIC

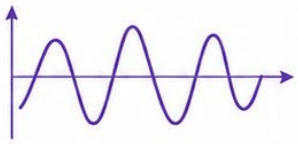

- frequency
- timing
- intensity

PEOPLE

- friendship
- familial association
- frequency of encounter

TIME

- elapsed or projected time
- past / present / future

EMOTION

- intensity
- immediacy

MATHEMATICAL / OTHER

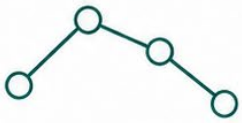

- dimensions appropriate to the information
- any applicable number of dimensions

2. AVAILABILITY ACROSS TIME

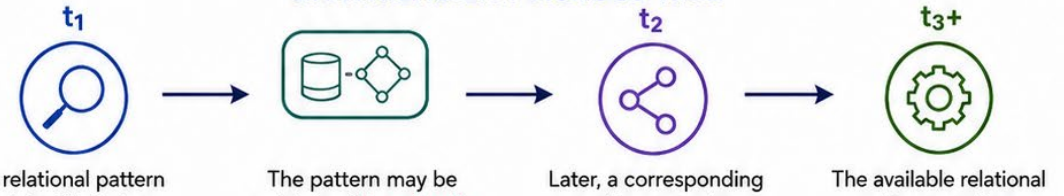

A relational pattern is identified at an earlier time.

The pattern may be maintained **transitorily** for a desired time or **persisted as a data entry.**

Later, a corresponding pattern may be encountered in the same or another space.

The available relational pattern may then support later processing.

Time

SPACES CAN OVERLAP

Time may be treated as a separate space or as a dimension of physical space.

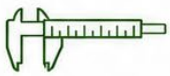

DIMENSIONS ARE SPACE-SPECIFIC

Relationships are evaluated using dimensions appropriate to the applicable information.

CONCEPTUAL DRAWING

Relative placement does not prescribe a required processing order, data structure, or physical location.

FIGURE 4B

The same F/E relational pattern across concepts and spaces

A persisted F/E, and a transitorily maintained F/E while it remains available, may support multiple concepts in the same or another space when their constituents exhibit the same relational pattern.

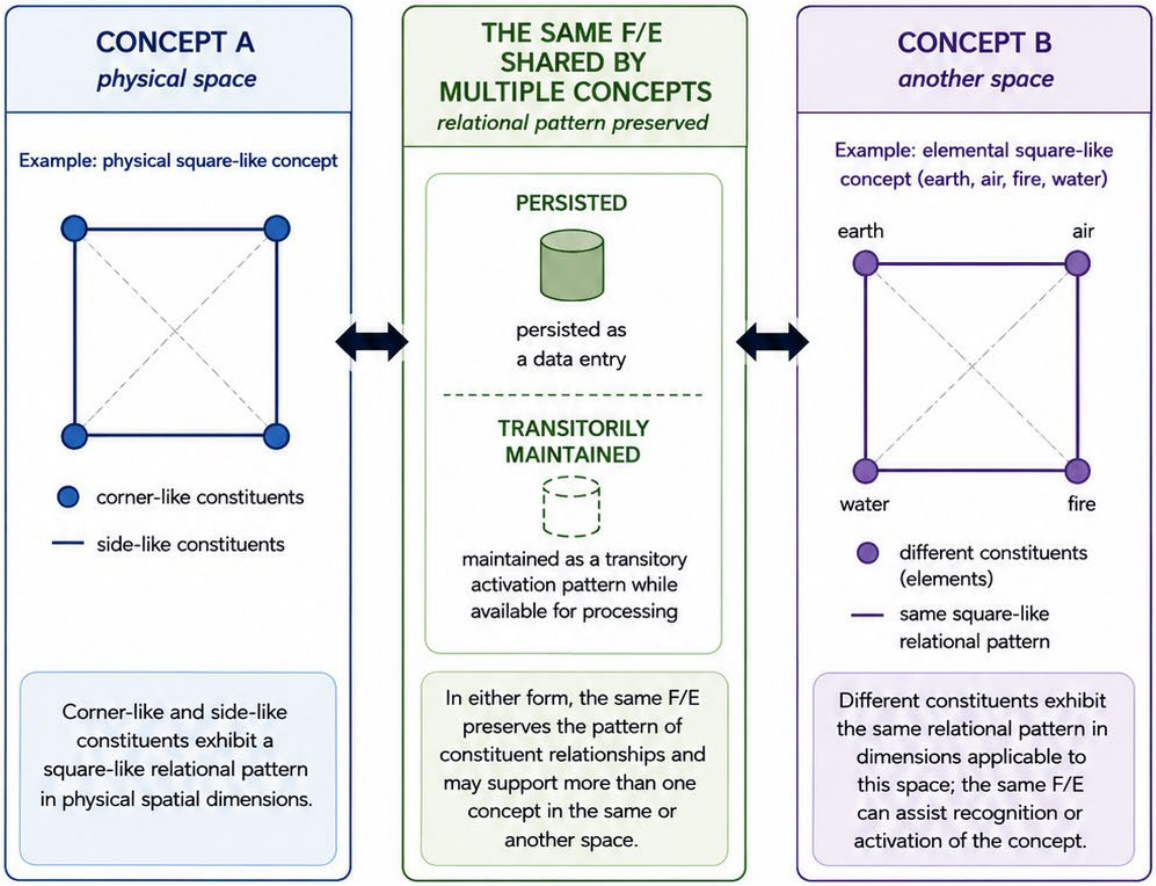

SUPPORTED COGNITIVE USES

- categorization
- identification / recognition
- association
- prediction
- activation / reconstruction
- higher-level pattern formation

These effects remain subject to the processed data, the active objective, applicable significance criteria, lifecycle treatment, and supervisory control.

REUSE DOES NOT CREATE IDENTITY

- Sharing the same F/E does not make the concepts identical.
- Their constituents remain different.
- Their meanings remain different.
- Their applicable dimensions or spaces may differ.
- The relative placement of the elements does not prescribe a required processing order or data-bank arrangement.
- No conversion of physical coordinates into nonphysical coordinates is implied.

The drawing is conceptual and illustrates reuse of the same relational pattern across concepts and spaces.

6. Significant patterns: criteria and recursive organization

TEAI uses significant-pattern processing to distinguish regularities warranting further treatment from the much larger number of possible arrangements in processed data. Processed data can originate outside TEAI or from current internal activity, including activation of concept or F/E data entries. In this overview, “significant” does not mean universally important, familiar, or already understandable to people. It means that the pattern satisfies one or more criteria applied to the relevant processing.

A pattern may be treated as significant because of a spatial, temporal, or other association among constituents; repetition or frequency of occurrence; a relationship to another significant pattern, event, or circumstance; unusual or distinguishable characteristics; a degree of order not characteristic of random information; an analogous occurrence at another level or in another space; or an empirical association with useful, provocative, unexpected, or otherwise informative results. These examples are not exhaustive. An implementation may use a restricted subset of such criteria, and the applicable criteria may change as experience shows that some are more useful than others.

Objectives and significance criteria play different roles. An objective can direct the system to seek a particular kind of pattern, select data sources or levels for examination, and prime concept or F/E data entries associated with the anticipated pattern. The significance criteria help determine whether an identified arrangement is treated as a significant pattern. Priming changes readiness to participate; it need not by itself establish significance or determine the pattern’s treatment.

Except where a data entry represents elemental data for the applicable operation, a concept or F/E represents a significant pattern defined by its constituents and their relationships. It can itself participate as a constituent of one or more further significant patterns. Those constituents can be persistent data entries or represented transitorily, and they can have the same recursive organization. Some processed data may be treated as elemental because TEAI does not further divide it by the means then available for that operation, not because no finer analysis is possible.

A fixed hierarchy of levels is not required. An implementation may organize data entries by pattern level, space, both, or neither. A higher-level pattern may draw on constituents from the same or different levels or spaces, and levels may be skipped when an intermediate persistent data entry is not useful. The important architectural relationship is between a significant pattern and its constituents, not a universal ladder through which every pattern must pass.

Identification as a significant pattern does not dictate persistence. A pattern may remain transitory, be persisted, contribute to a canonical form, or cease to be maintained. This selectivity allows recursive reuse without requiring every possible combination to become a persistent data entry.

Because data entries representing recursive patterns can remain linked to their constituent data entries, the architecture can preserve a navigable structural account of how a higher-level pattern was formed. That account can remain available for investigation even when the higher-level pattern does not yet have a familiar human meaning.

7. Concepts, filters/exemplars, and their attributes

Concepts and F/Es are two principal types of data constructs through which TEAI can represent significant patterns. The term “concept” is intentionally broad. A concept may concern a person, object, event, condition, idea, imagined possibility, or another pattern, whether concrete or abstract and whether or not people already understand or have named it. Words such as “house,” or “barn,” are useful labels for explanation, but the word is not the concept itself and need not appear in the data entry. In some implementations, however, a concept data entry may include an attribute linking it to a word associated with the concept.

An F/E generally preserves or applies a spatial relationship, or another relationship expressed through dimensions of an applicable space. In a filter-like use, it may pass, block, or note data satisfying a target condition; in an exemplar-like use, it may compare current data with a maintained reference pattern for categorization or identification. The slash joins two possible uses; an individual F/E need not perform both.

Concepts and F/Es have potentially overlapping roles. Concepts are generally useful for specifying what has been identified, while F/Es are especially useful for specifying the relationships through which constituents form a significant pattern. This distinction does not require separate software classes or an exclusive division of information. An implementation may combine the constructs, use only one type, or preserve in concept attributes relational information that another implementation preserves in an F/E.

When a concept or F/E is persisted, it is maintained as an identifiable data entry. In various implementations, the data entry can be examined separately and retain sufficient information to represent the concept or F/E. Its physical form may be a database record, table row, file, digital or analog hardware arrangement, or another suitable implementation. The architectural point is the identifiable representation, not a required storage technology.

A persisted concept or F/E data entry includes attributes that characterize it and affect how it participates. Linking attributes can associate the data entry with constituents of the significant pattern it represents, with higher-level patterns in which it participates, or with other related data entries. Links may be direct or indirect, and reciprocal linking is optional. A linking attribute may also

express a preference among alternative associations or indicate the degree or kind of deviation from a canonical form.

Other attributes may identify time of creation or expiration, importance or salience, confidence that the pattern satisfied applicable significance criteria, current activation intensity, or other useful information. Some implementations may associate optional behaviors with concept or F/E data entries. A behavior may specify or carry out a process, method, or function associated with the data entry; the operation may instead be performed by a separate functional element or another suitable software or hardware arrangement.

Except for a data entry representing elemental data for the applicable operation, a concept or F/E data entry represents its own significant pattern and can simultaneously participate as a constituent of several further significant patterns. The architecture therefore need not copy the entire data entry for every pattern in which it participates. Reuse through linking supports combinatorial range while retaining information needed to examine the formation and alteration of higher-level patterns. Data banks and functions may be combined, divided, or otherwise arranged; the central point is that selected concept and F/E data entries—including their attributes, linking attributes, and optional behaviors—remain identifiable so that their participation and alteration can be inspected and governed.

8. Participation, transitory patterns, and persistence

Participation and persistence describe different aspects of TEAI processing. Participation concerns the current salience, priority, relevance, or readiness of a concept or F/E, whether represented persistently or transitorily. Persistence concerns whether the concept or F/E is represented as a data entry. A persistent data entry can therefore participate strongly, weakly, or not at all in a particular operation.

The current participation of a persistent data entry may be expressed through an activation-intensity attribute. A relatively high value or another distinguishing state can indicate that the data entry is salient to current, potential, contingent, or predicted processing. Numerical values are not required; another technique can place the data entry in an elevated, intermediate, or reduced state. In this overview, “participation” is the broader functional idea, while activation intensity is one way of representing it.

An objective can change participation of a data entry without changing its persistence. A persistent data entry may remain represented in a data bank irrespective of whether its activation is increased or reduced. Priming does not by itself create the data entry anew, make it more persistent, change its persistent representation to a transitory one, or establish that the significant pattern it represents

is present in the processed data. It changes the data entry's current readiness or priority to participate.

A concept or F/E can be represented transitorily by maintaining a coordinated pattern of activation among the constituents that represent it. Those constituents can include persistent data entries and concepts or F/Es that are themselves represented by further transitory patterns. Transitory representation can therefore extend recursively without requiring a separate persistent data entry for every represented pattern.

"Transitory" does not necessarily mean brief. A transitory pattern can be maintained for any desired time or for the duration of specified conditions. Its activation and composition can change as processing continues, and it can participate in recognition, association, prediction, or a further significant pattern. A transitory concept or F/E can become persistent, and a persistent one can instead be represented transitorily.

"Persistent" likewise does not mean permanent or unchangeable. A persistent concept or F/E data entry can be changed, merged, pruned, transferred, or maintained only while specified conditions continue. Any affected linking attributes or higher-level patterns may require corresponding examination. Persistence is a form of representation and lifecycle treatment, not an irrevocable commitment.

Separating current participation from persistent and transitory representation supports both combinatorial range and directed correction. Many possible combinations can participate without each becoming a persistent data entry. If a data entry is receiving excessive present emphasis, its activation can be reduced without deleting it; if persistent representation is no longer warranted, the data entry can receive another lifecycle treatment. These operations remain subject to applicable objectives, core objectives, and the Enforcer.

9. Objectives and objective-guided priming

Objectives give current processing direction. In an illustrated implementation, an Objectives Manager determines the objective or objectives to be pursued and may formulate subobjectives useful in pursuing them. Such subobjectives may identify appropriate data sources, direct attention to selected portions of those sources, select a space or pattern level for analysis, account for relevant context, or initiate other processing. Multiple objectives and operations may be coordinated or carried out in parallel; the architecture does not require one fixed sequence.

In some implementations, an objective may itself be represented by a persistent or transitory concept or F/E. It may remain linked to the constituents that gave rise to it and may be associated with scripts, subobjectives, models, or other structures used in pursuing it. Its identity, attributes,

relationships, and current state can therefore participate in processing and remain available for authorized examination.

Objectives may be established or refined in response to internally generated activity or external processed data. Their attributes may indicate priority or importance; states such as dormant, active, on hold, or accomplished; and conditions, certainty, or probability relevant to determining completion. Because these features are represented within TEAI, objectives can be prioritized, monitored, revised, or concluded as governable parts of processing rather than treated as fixed commands applied from outside.

Objective-guided priming is one way an objective affects participation. A top-down process can place selected persistent data entries, or concepts or F/Es represented transitorily, in an elevated or preferential state. This can facilitate identification of corresponding significant patterns, make relevant constituents more ready to participate, or help maintain a transitory pattern through coordinated activation of its constituents. Priming changes current readiness or priority to participate. It does not by itself establish that an anticipated pattern is present, sufficiently supported, or appropriate for persistence or use. Priming can be selective across recursively related patterns or levels: an objective may include specified data entries, exclude others, or stop priming at a selected level. In the house-and-barn example, it can prepare shared constituents and distinguishing relationships without predetermining the interpretation.

Processed data can also affect participation. Identification of lower-level significant patterns may increase activation of linked concept or F/E data entries at the same, higher, or lower levels. Those changes may support higher-level patterns, while the active objective primes patterns or constituents from another direction. Objective-guided priming and processed-data effects may therefore converge on some of the same data entries. The objective affects readiness to participate; the processed data supplies information bearing on whether the pattern is supported. Neither direction alone need establish the result.

Figure 5 provides a conceptual view of this interaction using selected concepts as constituents. For visual simplicity, F/Es are not separately depicted, although they may also participate by preserving or applying relationships among constituents, including relationships shared by competing higher-level patterns. Bidirectional constituent connectors indicate that participation may be influenced through top-down priming or bottom-up processed-data effects; they do not imply identical operations in the two directions. The figure also shows predictive priming, examination of alternatives that may have been disadvantaged by top-down preparation, and several possible processing directions; it is illustrative rather than exhaustive and does not prescribe a required sequence, level structure, or allocation of functions.

FIGURE 5

Objective-guided and processed-data participation

Top–down priming and processed–data effects can meet across recursively related patterns.
Neither direction alone establishes which interpretation is sufficiently supported.

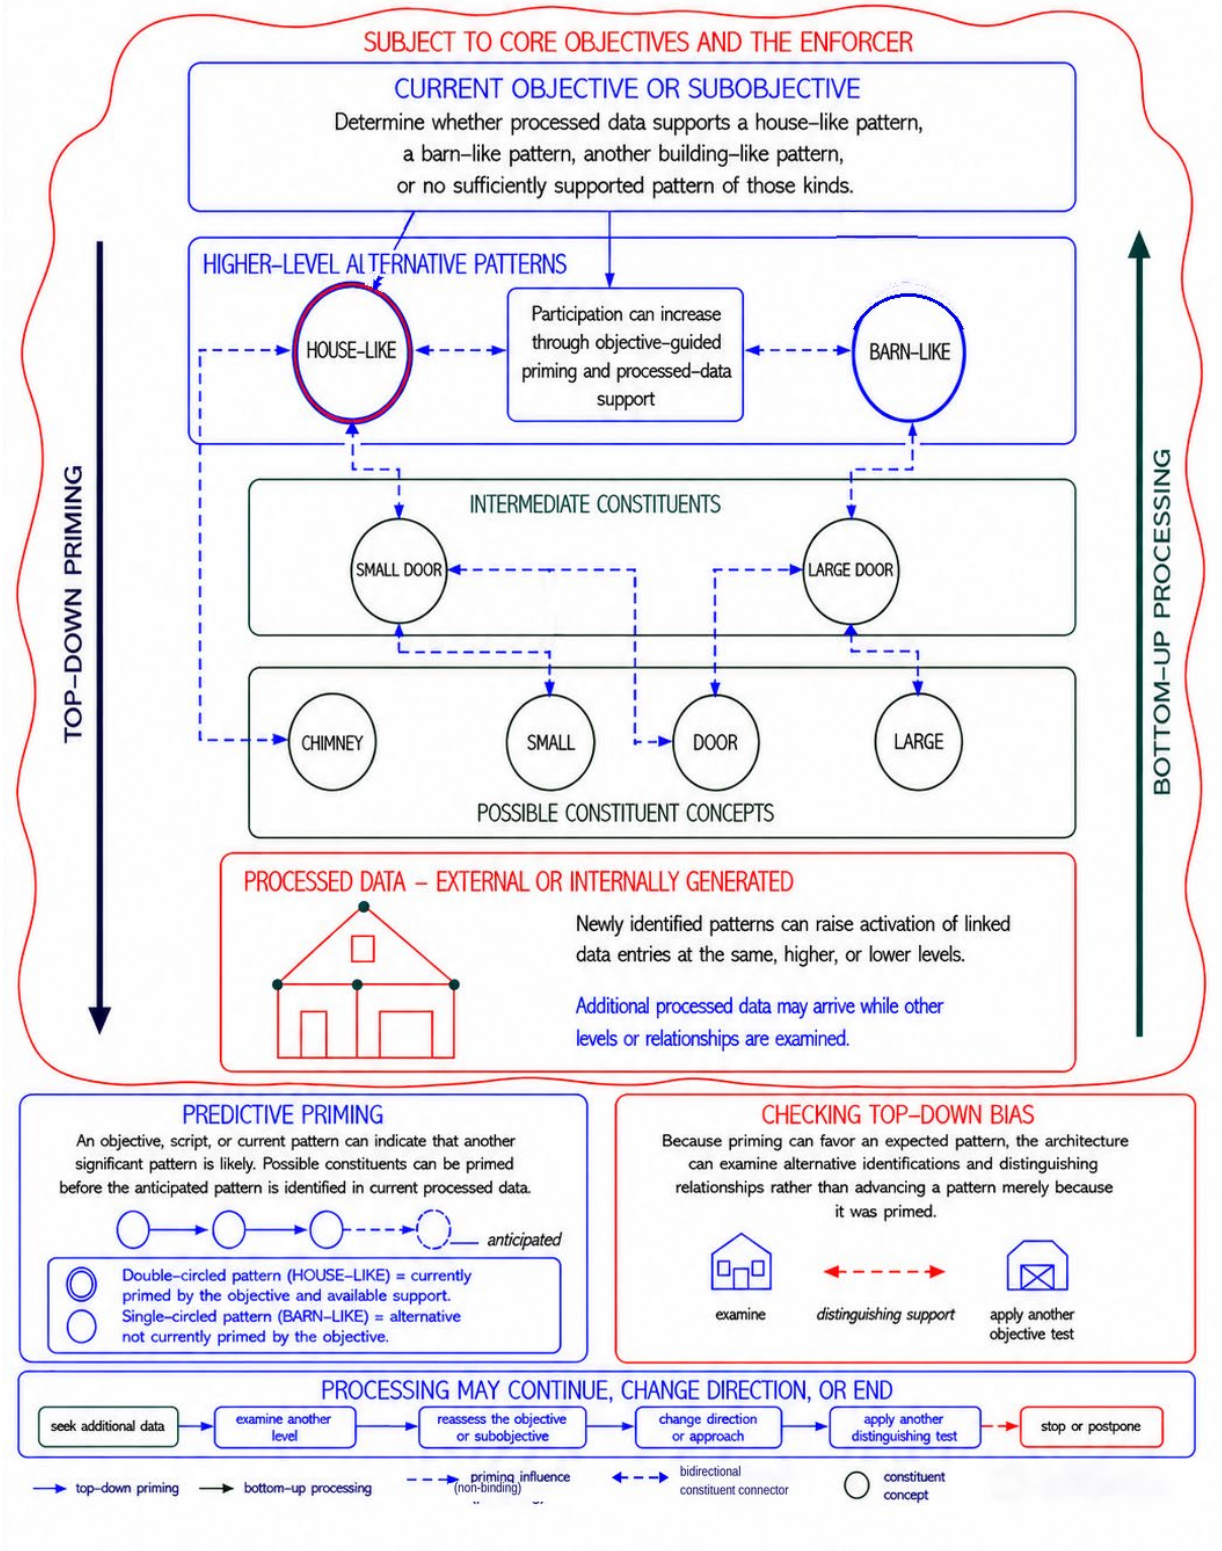

Priming can also be predictive. An objective, a script, or a currently identified significant pattern may indicate that another significant pattern is likely to occur. TEAI can then prime possible constituent data entries before the anticipated pattern is identified in the current processed data. In an illustrated implementation, a concept or F/E needed for such testing may be created with an attribute indicating that it is hypothetical and contingent on identification of the anticipated pattern. This preparation makes the expected constituents more ready to participate; it does not establish that the anticipated pattern is present or should be persisted.

Scripts can organize this predictive use of current information. A persistent or transitory script may link ordered or conditional sequences, probabilities, supporting factors, predicted constituents, and branches to other concepts or scripts. If current processing corresponds to one part of such a sequence, data entries associated with a likely later part can be primed. As processed data changes, another branch or sequence may become more relevant, and the priming can change accordingly. Because top-down priming can favor an expected interpretation, the architecture can examine alternatives and distinguishing patterns. In the house-and-barn example, this helps distinguish support arising from processed data from priority arising because a pattern was primed.

Objectives can also direct the search for additional data, examination of another level or space, or a change in processing direction. Any resulting lifecycle treatment remains dependent on the processed data, constituent relationships, applicable criteria, resource considerations, and supervisory control. Objective-guided processing remains subject to core objectives and the Enforcer. In an illustrated implementation, core objectives are established by authorized users external to TEAI and protected from alteration by the system’s other functional elements. The Enforcer can monitor objectives and related operations and can block processing of an objective determined to be contrary or injurious to a core objective. This supervision applies while objectives direct attention and priming, not merely after an output has been produced.

Because objectives and priming operate through identifiable data entries and activation states, their effects can remain available for investigation. Relevant information may identify the active objective or subobjective, the data entries that were primed, the significant patterns identified in processed data, the alternatives examined, changes in activation, and the persistent or transitory treatment applied. This can help an authorized reviewer distinguish what the processed data supported from what the objective made more ready to participate, without requiring reproduction of every lower-level machine event.

10. Alternative patterns, incomplete support, and correction

More than one significant pattern may remain active when the available processed data supports different interpretations. TEAI need not select one immediately. Alternative patterns can continue to

participate while the architecture examines their shared and distinguishing constituents, relationships, activation states, and support from processed data. In the house-and-barn example, many constituent patterns may participate in both interpretations, while distinguishing constituents and relationships, such as a chimney-like structure, window placement, opening size, scale, or roof configuration, may affect their relative participation.

Support may remain incomplete for several reasons. An expected constituent may be absent; a relevant relationship may not yet have been identified; the processed data may support more than one pattern; or the active objective may call for examination of another level, space, or source of data. Alternative patterns can remain transitory while these matters are considered. Concept or F/E data entries, scripts, and related processing may also include confidence, probability, or other information useful in evaluating support, but TEAI does not require uncertainty to be expressed through one numerical scale.

Comparison can occur at different recursive levels and in different forms. Higher-level patterns may share many constituents while differing in one or more relationships. Different F/Es may organize the same constituent data entries in different ways. A predicted pattern may be compared with a pattern identified in subsequently processed data. An observed arrangement may also be compared with a canonical concept or F/E, together with linking attributes indicating the degree or kind of deviation from that reference. By preserving the relevant constituents, relationships, and activation states, the architecture can compare alternatives without reducing each interpretation to an indivisible result.

Incomplete support need not be expressed only by reducing confidence in an entire interpretation. When identified constituents and relationships support only part of a larger significant pattern, TEAI may represent that subset as a partial concept, associate it with the larger concept or F/E it may help support, and preserve which expected constituents or relationships remain unresolved. Thus, two eye-like concepts and nose-related data may be associated both with a face-like concept and with a partial concept when they do not yet sufficiently support the complete face-like pattern. This allows the architecture to retain what has actually been identified without treating the larger pattern as established.

The Objectives Manager can direct further processing by seeking additional data, focusing on distinguishing patterns, examining another level or space, or reassessing the objective. These operations allow comparison to continue without treating an initially primed pattern as established.

Figure 6 summarizes this continuing comparison. It shows alternative patterns remaining available while the architecture investigates distinguishing support, directs additional processing, applies an appropriate treatment, or makes a targeted correction. The possible operations are not a required sequence, and the drawing does not imply that one alternative must ultimately be selected.

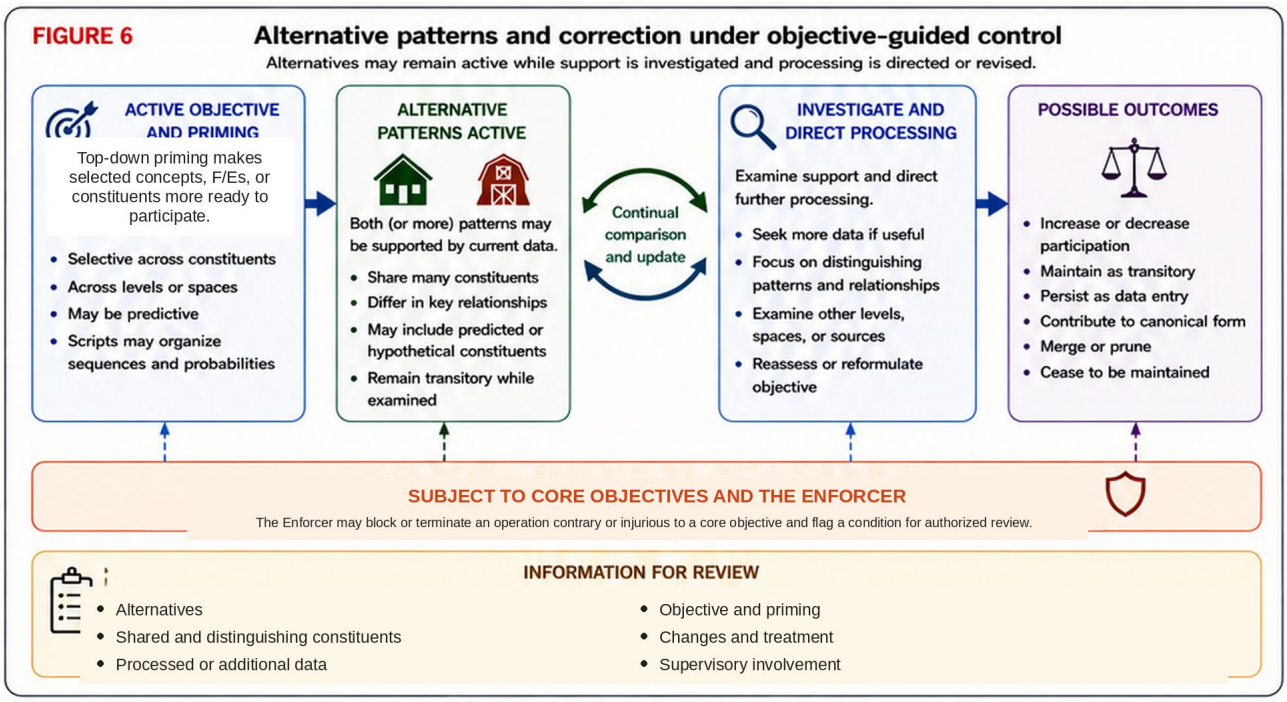

As processing continues, the relative participation of the alternatives may increase, decrease, or remain substantially unchanged. One or more patterns may remain transitory. A pattern may instead be persisted as a data entry, contribute to a canonical form, be merged or pruned, or cease to be maintained. Its activation may also be reduced if further processing supplies less support. The appropriate treatment depends on the processed data, constituent relationships, applicable significance criteria, current objectives, resource considerations, and supervisory control.

Identifiable data entries and relationships also permit more directed correction. If an objective is focusing processing too narrowly, it can be reassessed or reformulated without changing the processed data already received. If a concept or F/E data entry is receiving excessive emphasis, its activation-intensity attribute or other preferential status can be changed without deleting the data entry. If a linking attribute, another attribute of a concept data entry, or an F/E represents an overbroad or mistaken relationship, the applicable data entry can be examined and changed without discarding every associated concept. If a persisted pattern no longer warrants separate persistent representation, it can be merged, pruned, transferred, or represented transitorily under applicable procedures.

A directed correction need not be isolated in effect. Changing a linking attribute or F/E may alter higher-level significant patterns in which the affected relationship participates. Reducing activation may change which alternatives remain prominent. Merging or pruning a data entry may require examination and corresponding treatment of links involving that data entry. The architectural

advantage is not that every change is independent of the rest of the system, but that the structures directly affected by the change can be located and examined.

Corrective operations remain subject to core objectives and the Enforcer. Information available for review can identify the alternatives, distinguishing constituents and relationships, objective and priming, additional data, changes made, treatment applied, and any supervisory effect, at a level appropriate to the inquiry.

These arrangements do not establish that TEAI will always favor the better-supported pattern or make an appropriate correction. Processed data may remain incomplete, significance criteria may be poorly selected, an objective may direct attention unhelpfully, or a relationship may be represented inaccurately. An architectural benefit is that alternative interpretations, incomplete support, and corrective operations can remain identifiable and available for directed examination and governed change.

11. Canonical forms, deviation, and selective lifecycle management

Recognition often requires related variations to be treated as instances of a common pattern without erasing differences that matter. Face-like, triangle-like, house-like, emotional, and other significant patterns may occur in many forms. Persisting each variation as an unrelated concept or F/E data entry would increase the number of data entries that must be stored, searched, linked, and supervised. TEAI can instead create canonical concepts or F/Es that serve as references for related patterns.

A canonical form may be produced by analyzing a group of concepts or F/Es and identifying common features or relationships. Averaging may be used where appropriate, but the architecture does not require one method of forming the reference. The canonical form is itself a concept or F/E and can have constituents, attributes, linking attributes, optional behaviors, and persistent or transitory treatment like other such structures. It represents a maintained reference and need not reproduce any one observed instance.

Canonical forms can apply both to concepts identifying something and to F/Es preserving constituent relationships. A canonical face concept, for example, may represent common features of a group of face-like patterns. A canonical F/E may represent common spatial relationships among eye-like, nose-like, chin-like, or other constituents. A newly observed arrangement can then be compared with the maintained reference without requiring a preexisting data entry that precisely matches every feature.

Linking attributes to or from a canonical form can indicate the degree or kind of deviation between the reference and a related concept or F/E. The deviation may concern position, spacing, angular

orientation, scale, translation, intensity, or one or more other dimensions applicable to the relevant space. The architecture may also generate a measure of overall conformance between processed data and a significant pattern represented by the canonical form. The comparison can therefore preserve both the common reference and differences relevant to current processing.

Canonical organization can operate recursively. A canonical concept or F/E may represent a significant pattern as a whole, participate as a constituent of another significant pattern, or include canonical and noncanonical constituents. Those constituents may have the same range of forms. Use of a canonical reference therefore need not flatten recursive structure or replace the links through which constituent patterns remain identifiable.

Canonical forms can directly limit the number of separately maintained representations. A canonical reference, together with information indicating the degree or kind of deviation, may represent numerous related variations that otherwise could require separate persistent concept or F/E data entries or separately maintained transitory patterns. This does not eliminate the combinatorial expansion through which existing constituents participate in many higher-level patterns. It can limit the resulting proliferation of separately represented variations where shared structure and identifiable deviation provide an adequate representation.

The benefits may extend beyond storage. Fewer separately persisted data entries may reduce searching and coding demands. A maintained reference can make relevant deviations easier to detect, support comparison of processed data with an existing significant pattern, assist supervisory examination, and provide a clearer basis for presentation to authorized reviewers. The practical extent of these benefits depends on the implementation and the information being processed.

Canonical-form handling is one part of selective lifecycle management. In an illustrated implementation, the Persistence Manager can consider available resources, desired speed, recurrence, salience, and other operational factors when determining whether a concept or F/E should be created, changed, eliminated, represented persistently, maintained transitorily, or used in forming or modifying a canonical form. Usefulness in current processing does not by itself require creation of a separate persistent data entry. Merging and pruning are likewise selective operations. Merging can consolidate related persistent data entries or contribute to a more useful canonical form. Pruning can remove a data entry that is no longer useful or not justified in light of available resources and current needs. A concept or F/E data entry may remain linked both to its constituents and to higher-level patterns in which it participates. Accordingly, altering, merging, or pruning the entry may require examination of affected linking attributes and related patterns.

Figure 7 illustrates how observed variations can be compared with a canonical form and receive different treatment without requiring a separate representation for every variation.

FIGURE 7

Canonical forms, variation, and lifecycle discipline

A maintained reference can preserve shared structure while deviation information and treatment choices preserve meaningful differences.

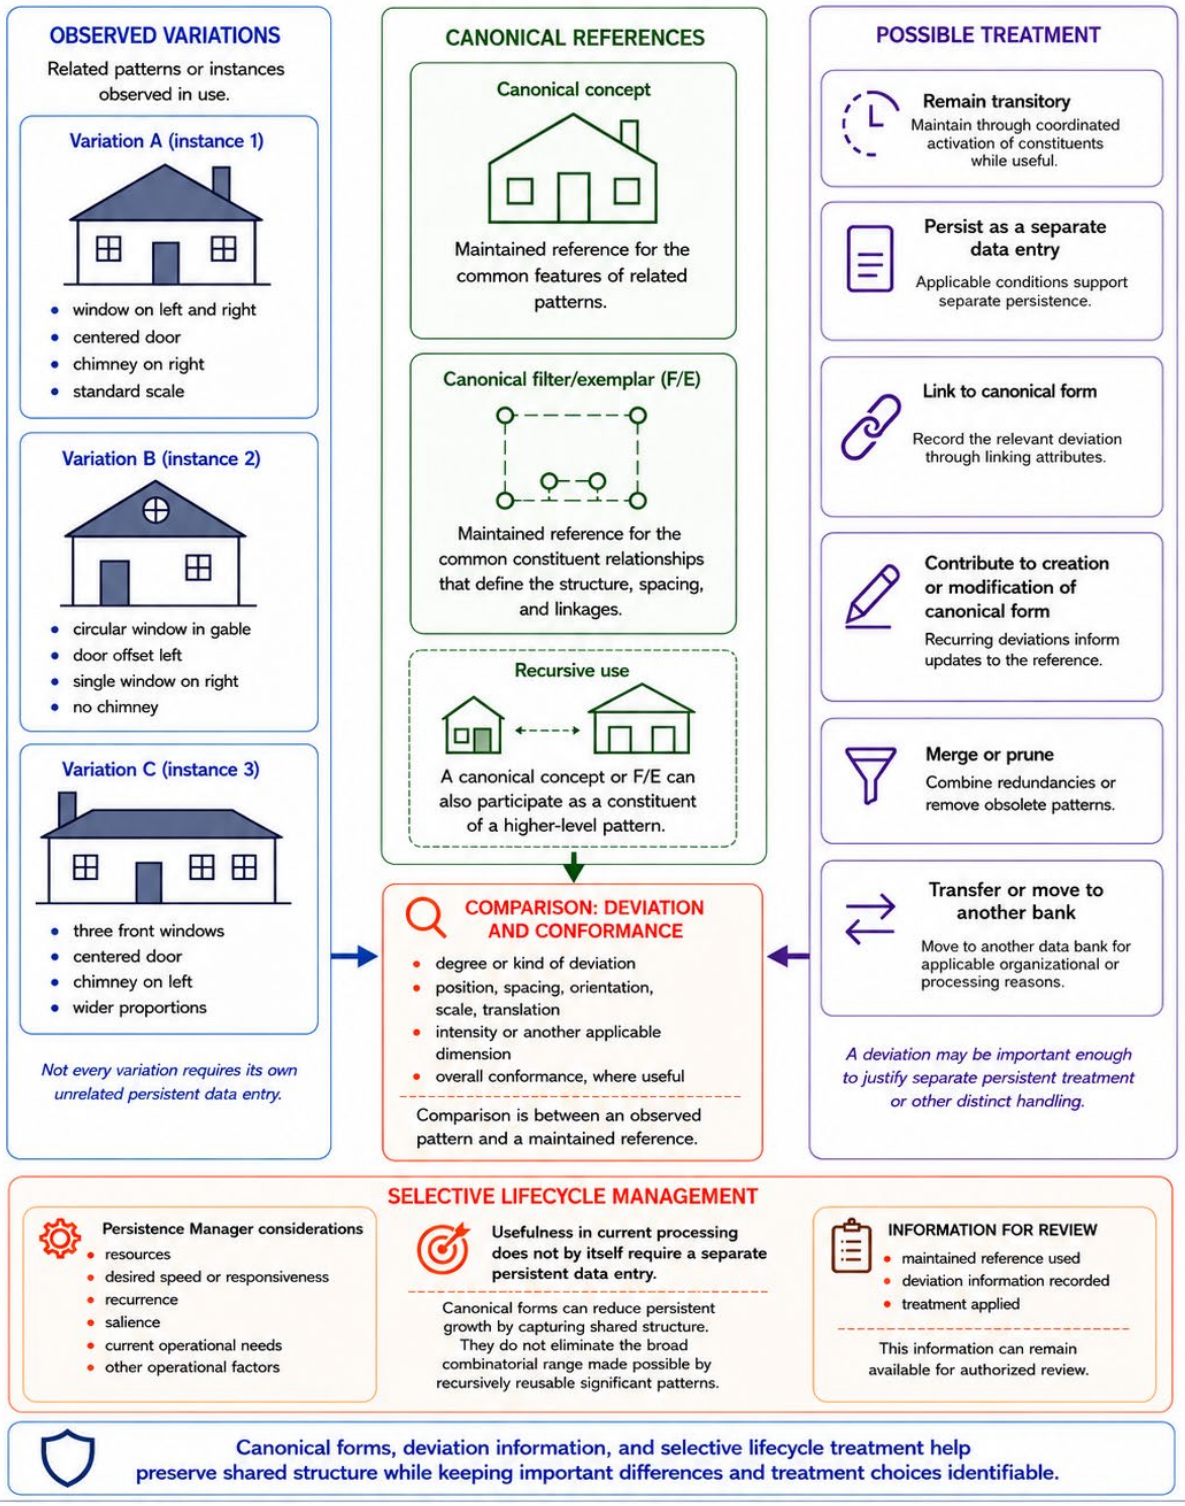

A canonical form does not require all deviations to be treated as equivalent or acceptable. A difference may be important enough that the observed pattern remains a separate alternative, receives its own persistent representation, contributes to modification of the canonical form, or receives another treatment. The architecture preserves the distinction between the reference and the deviation so that recognition across variation does not become unqualified generalization. Canonical forms, deviation information, selective persistence, transitory representation, merging, and pruning therefore operate together to preserve useful relational and pattern structure while limiting unnecessary representational growth. The result is not a requirement that all variations be compressed into one form, but an architecture in which shared structure and meaningful differences can remain separately identifiable and available for later processing and review.

Canonical forms can also provide an operational reference for determining whether an observed arrangement sufficiently corresponds to an existing significant pattern. In an illustrated implementation, comparison may consider deviation in one dimension or in a combination of dimensions, and identification of the pattern may depend on whether those deviations remain within ranges appropriate to the pattern and the information being processed. This allows recognition across variation without requiring every difference to be ignored or every variant to receive a separate persistent representation. Because the canonical reference, the relevant deviations, and the resulting identification remain represented through identifiable structures and relationships, the basis for treating an observed variation as corresponding — or not sufficiently corresponding — to the reference can remain available for examination.

12. Governance, supervision, and auditability

Governance in TEAI concerns how authorized requirements can affect objectives, participation, change, persistence, use, external action, and other aspects of processing. It can operate throughout processing. A significant pattern may be identified yet remain transitory. A concept or F/E data entry may be subject to authorization before it is changed. Information from an external source may require further processing before it affects persistent structure or external action. Output processing and transmission remain subject to applicable core objectives and supervisory control.

In the illustrated architecture, core objectives are established by authorized users external to TEAI and protected from alteration by the system’s other functional elements. These core objectives may embody requirements supporting alignment with human-established values and interests. The Enforcer monitors objectives and data entries for individual and collective consistency with the core objectives and may block processing of an objective determined to be contrary or injurious to one of them or take other remedial action. Supervisory functions may be combined with or distributed among other functional elements without changing their governing role.

Governance also includes controlled access and alteration. Known or future security techniques may restrict who or what can provide data, obtain output, inspect the system, or alter its operation. Access may be restricted to particular users or groups, and authorization may govern such interaction. Transparency does not require unrestricted disclosure or universal authority to change the system.

Auditability, as used here, concerns making information relevant to an operation available for inspection or reporting. Depending on the implementation and purpose, the available information may identify, for example, the objective or subobjective involved; the concept or F/E data entries and constituent relationships that materially participated; alternative significant patterns that remained under consideration; the treatment of external-source information; changes in activation or persistent or transitory representation; supervisory or authorization actions; and the resulting output or other external action. A concise explanation may identify the principal objective, participating patterns, and result. A more detailed view may show constituent data entries and relationships. A deeper review may examine changes in activation, linking attributes, persistent or transitory treatment, and supervisory action.

Figure 8 brings these relationships together, showing core objectives and the Enforcer acting across objectives, priming, significant-pattern processing, lifecycle treatment, authorized change, and external action. It also shows that review information may be presented at different levels without implying one fixed processing sequence or unrestricted access.

Governance may operate at more than one stage of processing rather than only at the point of final output. For example, supervision may concern whether an objective should proceed, whether additional data or further processing should be undertaken, whether a significant pattern should remain transitory or be persisted, or whether another lifecycle or supervisory action is appropriate. This allows supervision to operate during processing rather than only as an after-the-fact check.

Supervision may also be prospective. The Enforcer can monitor processing and, in some implementations, assess whether an objective or anticipated action may be contrary or injurious to a core objective, including through predictive scripts. If a concern is identified or predicted, the Enforcer may block processing or take other remedial action.

FIGURE 8 Governance, supervision, and review surfaces

Authorized requirements can affect objectives, participation, lifecycle treatment, change, and external action, while review surfaces provide different levels of explanation for different authorized roles.

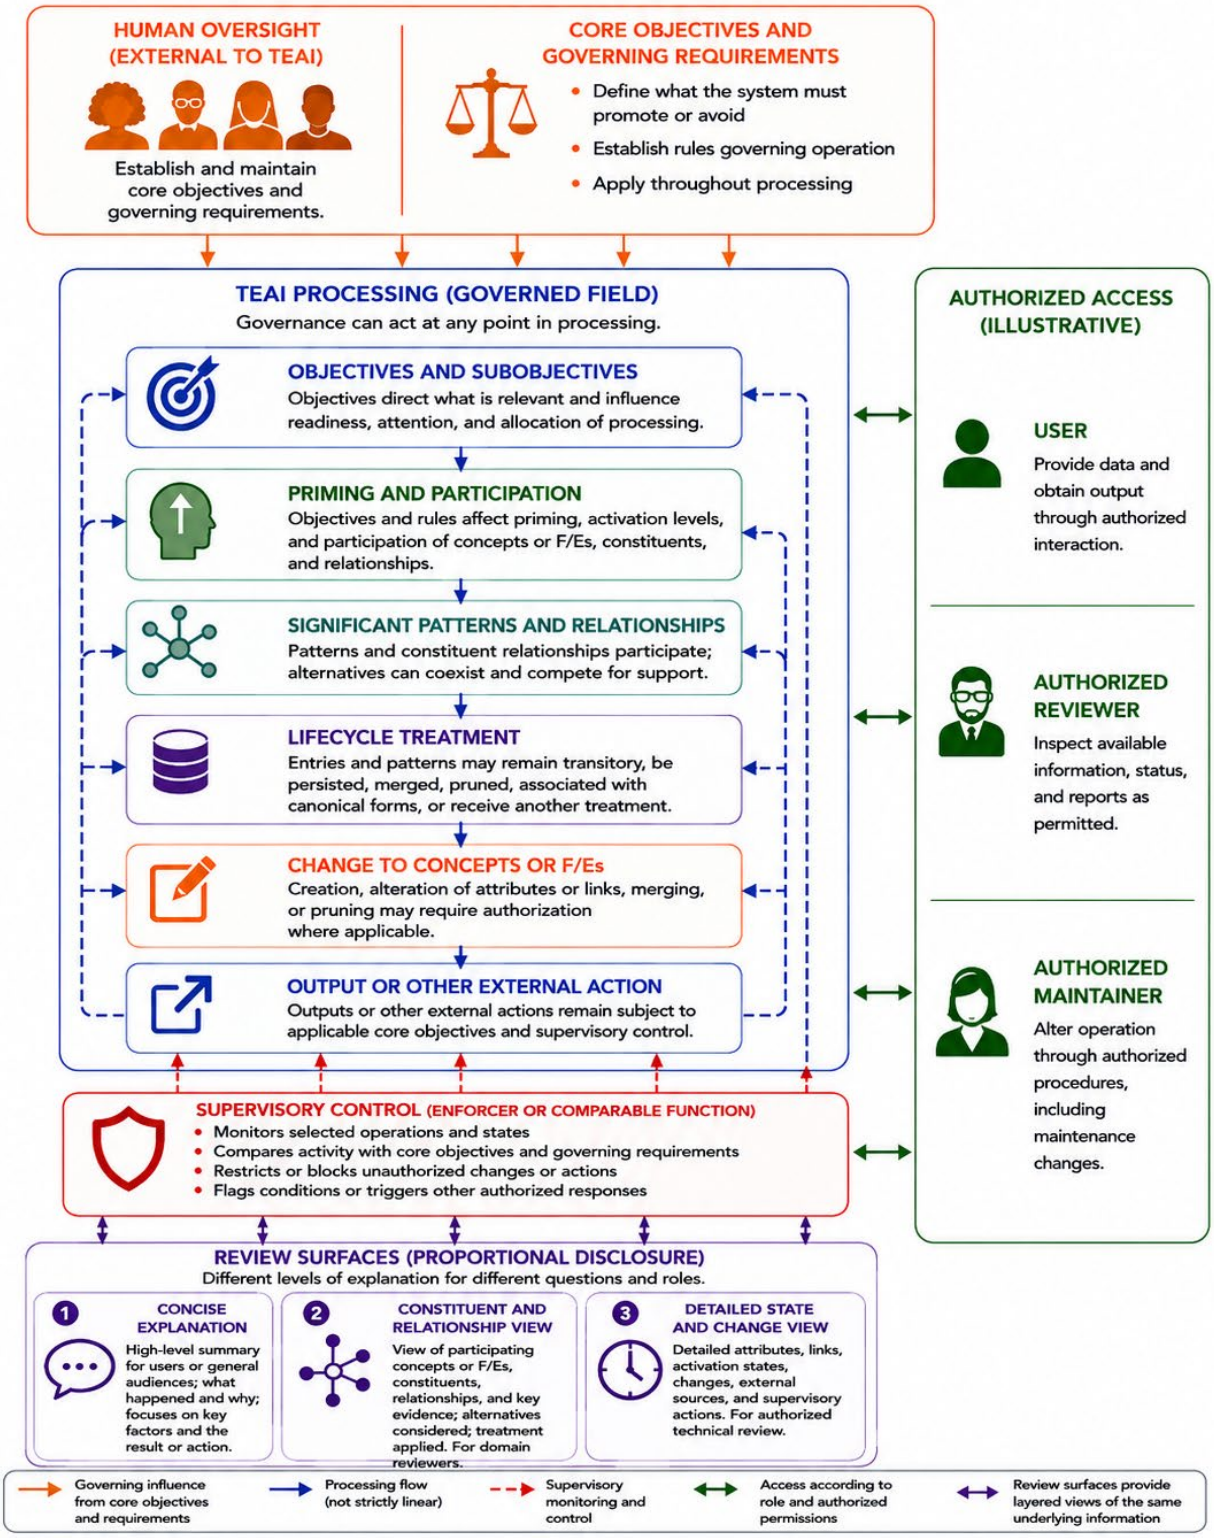

Review should preserve alternatives and incomplete support when they are relevant. A presentation showing only the final treatment can create a misleading impression that the result was inevitable. The available information may therefore identify alternative significant patterns, missing constituents or relationships, unresolved conditions, and the effect of objective-guided priming. The purpose is to make processing examinable, not to convert a qualified result into an oversimplified narrative.

Security and privacy remain part of governance. Inspectable data entries and review information may contain sensitive or proprietary material that should be available only to authorized users. Practical implementations may therefore restrict access, require authentication, control alteration, and limit the information presented for a particular review. “Transparent” need not mean publicly or universally accessible.

Governance mechanisms themselves require examination. A poorly formulated core objective may constrain the wrong operation. An incomplete review may omit an important participating structure. An overbroad authorization may weaken encapsulation. A supervisory operation may introduce its own failure modes. TEAI provides identifiable structures and operations through which these matters can be examined and controlled, but implementation quality and human judgment remain important.

These features may be particularly useful where a result must be reviewed, corrected, or produced under authorized requirements. Their value should be evaluated in the applicable implementation and setting. Inspectable information does not establish that an objective was sound, that a governing requirement was justified, or that a result was fair. Rather, it supplies a more concrete basis for examining those questions.

Representative review questions

What objective or subobjective was active, and who or what authorized it?

Which concept or F/E data entries, constituents, and relationships materially participated?

What alternative significant patterns remained under consideration, and what support was incomplete?

What significance criteria caused a pattern to be treated as significant, and what supporting information satisfied those criteria?

What effect did objective-guided priming have on which data entries, constituents, or relationships were more ready to participate?

What predicted, hypothetical, or contingent patterns were considered, and what later information supported or failed to support them?

Was a canonical concept or F/E used as a reference, and what deviation or conformance information affected the comparison?

What remained transitory, what was persisted, and what was changed, merged, pruned, or otherwise treated?

What external-source information contributed, and how was it treated?

What core objective, Enforcer action, authorization, review, or other supervisory action affected the result?

13. External data and complementary components

TEAI may receive data or proposed results from many sources, including sensors, databases, user interfaces, network services, statistical models, neural networks, transformer-based systems, specialized processors, and other tools. A source may be physically separate from TEAI or may operate within a common computing system while remaining external to TEAI's architecture. In either case, its output can enter TEAI as external input and be processed without requiring TEAI either to reproduce the source's internal methods or to treat the output as an established internal conclusion.

The principal architectural question is how TEAI handles the material after receiving it, regardless of whether the source is physically separate or operates within a common computing system. Depending on the implementation and purpose, the material may be formatted, filtered, cleaned, verified, error-corrected, combined with other data, or otherwise processed, and TEAI may preserve information identifying its source or form, timing, reliability, noise, confidence, or relevant relationships. The resulting processed data may be compared with existing concept or F/E data entries, maintained only for current processing, persisted as a data entry, associated with an existing data entry, or receive no further treatment.

A learned component may therefore contribute useful capabilities without its internal representations becoming part of TEAI's addressable structure. A transformer-based system may generate or interpret language; a vision model may identify features or propose a classification; a database may provide recorded information; and a sensor may supply physical measurements. TEAI can use such material as processed data, as a basis for further significant-pattern processing, or as one contribution among several. The complementary component's internal process remains distinct from the concept or F/E data entries—including their attributes, linking attributes, and activation states—through which TEAI uses the resulting information, whether the component is physically separate from TEAI or operates within a common computing system.

This arrangement permits complementary use without assuming that the external output is complete or authoritative. External information may be incorrect, stale, biased, noisy, insecure, unavailable, or unsuitable for the current objective. TEAI may therefore restrict it to current processing, compare it with other information, seek additional data, prevent it from automatically changing protected data entries, or leave the resulting processing subject to the same supervisory control as other TEAI operations.

Information available for review may identify the source or form of the external material, the role it played, whether it was transformed or compared with other information, whether it affected a persistent data entry or remained transitory, and whether a core objective or the Enforcer affected its

use. This preserves an inspectable boundary between complementary components and TEAI processing without requiring every contributing system to share TEAI's internal architecture.

14. Practical evaluation and architectural significance

Identifiable structure brings engineering demands as well as opportunities. TEAI is organized around identifiable significant patterns, recursive constituent relationships, objectives, and governed lifecycle operations. This organization may be especially suited to modeling or simulating selected patterns of human cognition and behavior, including exploring how people or groups may interpret or respond to information, examining thought patterns and biases, interpreting communications, and supporting clearer understanding between people or groups by showing how communications may be received by different audiences. Engineering considerations may include storage, organization, retrieval, selective alteration, lifecycle treatment, and other operations involving concepts, F/Es, and their relationships. Recursive processing can support a large range of combinations, while review information may require selective organization and presentation. Interfaces must answer the relevant question without obscuring it in unnecessary detail. Practical implementations must balance expressive range, processing cost, representational growth, review quality, and ease of use.

Data quality, significance criteria, objectives, and supervisory requirements require testing and revision. Identifiable structure can make a problem easier to locate, but it does not prevent incorrect information, poorly chosen criteria, mistaken relationships, or unhelpful objectives from affecting processing. Core objectives and the Enforcer likewise require appropriate formulation, authorization, protection, and evaluation. Engram Factory is advancing selected parts of the framework toward working implementations and bounded demonstrations for evaluation under intended operating conditions.

TEAI addresses a concrete architectural question: can substantial cognitive processing be carried through internal structures whose identity, constituent relationships, changing participation, treatment, and governing objectives remain available to the machine and to authorized human review? It approaches that question by bringing significant-pattern processing and human-directed governance into the same operational organization. Its proposition is not that inspectability alone produces correctness or safety, but that recognition, association, prediction, reasoning, learning, persistence, correction, and supervisory control can operate through the same identifiable internal structures and relationships.

TEAI therefore offers a practical architectural path toward useful cognitive capability in systems whose important internal structures and changes are more understandable, selectively correctable, and subject to human-established control. The extent of that potential will depend on implementation and evidence, but the architecture provides a disciplined basis for pursuing it.