Foundations of General Systems Theory from a Critical Realist Perspective

John A Challoner, July, 2026

Abstract

General Systems Theory has developed through contributions from many disciplines, often resulting in differing terminologies and overlapping concepts. This paper reconstructs the conceptual foundations of General Systems Theory from a Critical Realist perspective by distinguishing between the organisation of reality (ontology) and the representations used to describe it (epistemology).

Beginning with the concepts of entities, properties, relationships, interactions and processes, the paper develops a coherent progression through collections, configurations, assemblies and systems, showing how increasingly organised causal interactions give rise to emergent properties, behaviours and outputs. Throughout, concepts are defined in plain language while maintaining conceptual precision and distinguishing carefully between physical phenomena and their analytical representations.

The paper further demonstrates that systems concepts apply recursively across successive levels of organisation, providing a common conceptual framework for describing physical, biological, engineered and social systems. This recursive organisation explains how systems become components of larger systems while preserving a consistent systems vocabulary across levels of analysis.

Rather than proposing a new systems theory, the paper clarifies and organises the conceptual foundations underlying existing General Systems Theory. It provides a coherent systems idiom intended to improve communication across disciplines and to establish the conceptual basis for subsequent papers examining causality, emergence, constraints, feedback, and systems dynamics.

1. Introduction

Systems theory is concerned with the organisation of reality and with the ways in which organised entities interact, change and persist. Before introducing foundational systems concepts, however, it is necessary to clarify the relationship between reality and our representation of it.

This paper adopts the Critical Realist position that reality exists independently of our knowledge or description of it (Bhaskar, 1975; Bhaskar, 1978). Human beings do not create reality by observing, classifying or naming it. Nevertheless, our knowledge of reality is necessarily mediated by cognition. The entities, relationships, processes and organisations that we recognise are represented through cognitive and linguistic structures that have developed within humanity.

As discussed in *Systems Linguistics: Representation, Compression, and Productive Coordination* (Challoner, 2026), these capacities have evolutionary foundations. Organisms capable of recognising configurations, changes and causal regularities in their environments are better able to identify threats, exploit opportunities and act in ways that support their continued viability (Friston, 2010). Cognitive representations need not be complete or perfectly accurate to provide such advantages. However, representations that bore no useful relationship to the causal organisation of reality would be unlikely to support effective action and persistence.

There are therefore good reasons to suppose that human cognition preserves at least some causally significant features of reality. Our representations are neither reality itself nor arbitrary constructions. They are evolved, compressed and imperfect representations that are sufficiently related to reality to permit, within limits, successful explanation, prediction and intervention.

Systems theory builds upon these basic cognitive capacities. We distinguish entities, recognise relationships and interactions between them, observe processes and changes through time, and identify relatively stable organisations. We may then recognise an organisation as a whole and represent it as a system.

A distinction must therefore be maintained between an organisation that exists in reality and its representation as a system. The representation is epistemic: it is a means by which an observer identifies, simplifies and investigates an ontological organisation (Bhaskar, 1975). The two are not identical. Boundaries may be drawn differently, different components may be emphasised, and different relationships may be included according to the purpose and perspective of the investigation. Nevertheless, where a systems representation preserves causally significant features of the organisation it represents, it can provide a useful basis for understanding and action.

The purpose of this paper is to establish a foundational set of systems concepts in plain language. Beginning with entities, relationships, interactions and processes, it develops the concepts of organisation and system before considering components, boundaries, environments, inputs, outputs, states, behaviour and nested systems. These concepts provide the foundations for the subsequent consideration of causality and emergence.

2. Entities and Their Characteristics

2.1 Entities

Reality is complex and continuous. To understand and reason about it, we distinguish portions of reality from their surroundings and treat them as identifiable things. In this paper, such identifiable portions of reality are referred to as entities.

Many entities are readily recognised. A stone, a tree, a person, and a planet may each be identified as an entity. However, the concept is not restricted to familiar material objects. A field, a relationship, or an event may also be treated as an entity where it forms an identifiable portion of physical reality. The concept is therefore deliberately broad.

The ability to identify entities depends upon human cognitive processes. We encounter recurring similarities within reality and learn to recognise configurations that share common features (Rosch, 1973; Rosch, 1978). A tree may differ considerably from every other tree in its precise shape, size, colour, and structure, yet recurring similarities allow different trees to be recognised as instances of the same general kind of entity.

For familiar concrete entities, these similarities may be compressed into a simplified iconic representation. Rather than retaining the complete configuration of every tree encountered, cognition retains those recurring features that are particularly useful for recognition. The resulting representation may be understood as a prototype (Rosch, 1978): a simplified representation of the common or typical features associated with a category of entities.

The prototype is not necessarily a detailed internal picture. It is a cognitively compressed representation containing sufficient recurring features to assist recognition. When a new entity is encountered, its perceived features may be compared with this representation and it may consequently be recognised as, for example, a tree, dog, chair, or person.

Human symbolic representation adds a further layer of compression (Bruner, 1966). The iconic or conceptual representation of a tree may be associated with the linguistic symbol *tree*. This allows the recognised category to be communicated to others without reproducing the underlying perceptual representation. Iconic and symbolic representations are therefore related but distinct: one assists recognition, while the other greatly extends communication and reasoning.

Some entities cannot readily be represented through a single iconic prototype. Relationships, events, and concepts such as trust, competition, democracy, or causality do not possess a characteristic physical appearance that can be represented in the same way as a tree or a stone. These are referred to here as abstract entities.

Abstract entities may nevertheless be recognised through recurring similarities. Instead of relying primarily upon an iconic prototype, cognition may retain exemplars: remembered or represented instances of situations in which the abstract entity is understood to pertain. A person may recognise competition by relating a present situation to familiar examples such as a race, rival businesses seeking customers, or animals competing for resources. The examples differ configurationally, but are understood as instances of a recurring underlying relationship or process.

Humans also construct iconic analogues that provide visual representations of abstract concepts. Mathematical graphs, state-space diagrams, causal networks, maps, and similar representations do not depict the abstract entities themselves but serve as visual analogues that support reasoning about them. Such iconic analogues play an important role in the comprehension and communication of abstract concepts.

This distinction between the representation of concrete and abstract entities should not be regarded as absolute. Concrete entities may be recognised through exemplars as well as prototypes, and abstract entities may acquire conceptual prototypes. The distinction is instead one of representational emphasis. Concrete entities frequently permit configurational similarities to be compressed into iconic prototypes. Abstract entities depend more heavily upon symbolic representation, conceptual compression, and recognition across multiple exemplars.

The identification of an entity therefore contains an epistemic element. Different observers may recognise entities at different scales, classify them differently, or associate an abstract concept with different exemplars. Nevertheless, within the framework developed here, entities are not regarded as purely conceptual constructions. The act of identification is epistemic, but that which is identified is assumed to form part of physical reality (Bhaskar, 1975).

The definition adopted in this framework is therefore:

Entity: Any identifiable portion of physical reality, including objects, fields, relationships, and events.

This broad concept of an entity provides the first building block for systems theory. Entities may possess characteristics, enter into relationships, undergo change, and form organised structures. These concepts are considered in the following sections.

2.2 Properties and Characteristics

Once an entity has been identified, it may be described in terms of its properties or characteristics. These are features used to describe, distinguish, or classify entities. A stone may have a particular mass, shape, colour, and chemical composition. A tree may have a particular height, age, species, and branching structure.

In ordinary language, the terms *property* and *characteristic* are often used interchangeably. The same practice is adopted here. A property or characteristic is a recurring feature used to classify or distinguish entities.

Properties are often themselves abstract entities. Mass, height, age, colour, and temperature are not normally apprehended as separate concrete objects. They are recognised as recurring features of physical entities and represented through abstract concepts. A particular entity may possess a particular value of a property, while the property itself may recur across many different entities.

An important consequence follows from this interpretation. Where one or more recurring properties are used to identify an abstract entity, the entities possessing those properties may be treated as components of that abstract entity. Justice is an abstract property. A court judgment that is recognised as just is a concrete instantiation of that property. The concept *justice* may be regarded as a component of the more inclusive concept *virtue*, while the concept *virtue* may itself be regarded as a component of the broader concept *ethical quality*.

This introduces a form of nested hierarchy. The same entity may simultaneously form part of several progressively more inclusive abstract entities according to the properties used to identify and aggregate them. A particular court judgment may be represented as a just judgment, a virtuous judgment, or a judgment exhibiting ethical quality. The judgment itself has not changed. What changes is the level of abstraction at which it is represented.

This aggregational hierarchy should be distinguished from the organisational hierarchy considered later in this paper. In an organisational hierarchy, components interact causally to form a larger organised entity. In an aggregational hierarchy, entities are treated as components of a larger abstract entity because they share recurring properties. Both involve the representation of multiple entities as a larger entity, but the basis of their aggregation differs.

The relationship between properties and abstract entities also helps to explain the different forms of cognitive representation discussed in the preceding section. A prototype (Rosch, 1978) provides a compressed representation of the recurring properties that distinguish an abstract entity. An exemplar, by contrast, represents a particular or typical component through which the abstract entity is instantiated. The prototype assists recognition through shared characteristics, while exemplars assist recognition through comparison with previously encountered instances. For highly abstract concepts, however, prototypes and exemplars may be insufficient to communicate the relationships or organisation involved. In such cases, humans often construct iconic analogues, such as graphs, diagrams, or maps, which provide visual representations that support reasoning about abstract entities and their relationships.

The preceding discussion has focused upon the ontological relationship between properties and abstract entities. Before continuing, it is useful to briefly consider how this relationship influences cognition and communication.

The linguistic symbol (Bruner, 1966) associated with an abstract entity provides a further form of representation. The symbol *justice*, for example, may become associated with prototypical representations of the recurring distinguishing properties associated with justice, remembered exemplars of just actions or judgments, and iconic analogues that assist reasoning about the concept. The linguistic symbol allows the abstract entity to be communicated without reproducing these underlying cognitive representations, relying instead upon the recipient to reconstruct an appropriate understanding from their own prior knowledge and experience.

The linguistic symbol itself does not contain meaning (Saussure, 1916/2011). Rather, it acts as a trigger that activates cognitive representations already held by the recipient. Understanding therefore depends not only upon the words used, but also upon the extent to which the speaker and recipient share compatible prototypes, exemplars, iconic analogues, and other cognitive representations associated with those words. Communication is therefore not the transfer of meaning itself, but the transfer of symbolic representations from which meaning is reconstructed.

Having considered their cognitive and linguistic representation, we may now return to the ontological role of properties in defining, recognising, and organising abstract entities.

The three forms of representation may be distinguished as follows:

A prototype represents recurring distinguishing properties. An exemplar represents a particular or typical instantiation. A symbol provides a communicable representation of the entity.

These forms of representation are complementary rather than mutually exclusive. Concrete entities may be recognised through comparison with exemplars as well as prototypes, and abstract entities may be represented through combinations of conceptual prototypes and remembered instances. The relative importance of each form of representation depends upon the entity being recognised and the cognitive resources available to the observer.

Properties are also important because they enable comparison. Two entities may differ in many respects while sharing one or more properties. Stones of different shapes may have similar chemical compositions; different animals may possess similar body temperatures; and systems of very different kinds may exhibit stability or resilience. Recognition of recurring properties allows similarities and differences between entities to be represented without treating the entities themselves as identical.

Not all properties are immediately observable. Some require measurement, investigation, or inference. The mass of an object may be measured, while the internal structure of a distant star may be inferred from observable effects. The inability to apprehend a property directly does not imply that the property is unreal. As with abstract entities more generally, the distinction between reality and our means of recognising and representing it must be maintained.

The definition adopted in this framework is:

Characteristic (Property): A feature common to entities used to classify or distinguish them.

Properties provide the basis for describing the state of an entity. Before considering state, however, it is useful to distinguish properties capable of taking different values.

2.3 Continua, Dimensions and Measurement

Many properties of entities may be represented by assigning them values on one or more dimensions. Within this framework it is useful to distinguish between the physical continua and properties that exist independently of observation and the representational systems developed to describe them.

Physical reality includes continuous aspects such as space and time. Entities occupy positions and extend through these continua, allowing them to be represented as points, lines, areas, volumes, instants, or intervals. The continua themselves are assumed to exist independently of observation. However, the coordinate systems used to represent them, such as Cartesian coordinates (Descartes, 1637), latitude and longitude, calendars, and clocks, are epistemic constructs developed to identify, compare, and communicate positions and intervals according to the purpose of an investigation.

Many other properties, including mass, temperature, pressure, probability, and income, are likewise represented using dimensions. In these cases, the underlying property is assumed to exist independently of observation, while the measurement dimension, units, scales, and numbering system used to represent it are symbolic constructs developed for communication and analysis.

Representation therefore involves a sequence of mappings from reality to progressively more compressed symbolic forms. Consider an asteroid. The asteroid occupies a three-dimensional region of physical space. If the investigator is interested only in its length, the spatial span corresponding to that length is selected from the asteroid's overall geometry. That span is then mapped onto a point on a length dimension, and finally onto a numerical value using an agreed unit of measurement such as metres or kilometres.

Each stage of this sequence compresses information while preserving those aspects relevant to the analytical purpose. The complete geometry of the asteroid is progressively represented by a single property, then by a position on a measurement dimension, and finally by a numerical value in the number dimension. Measurement should therefore be understood not simply as assigning numbers to reality, but as a sequence of purposeful representational mappings.

This relationship between reality and representation reflects a general principle developed throughout this framework. Physical reality provides the continua and properties that exist independently of observation. Observers develop coordinate systems, dimensions, scales, units, and numbering systems to represent those continua and properties in forms suitable for communication, comparison, and reasoning.

The definitions adopted in this framework are therefore:

Continuum: A continuous aspect of physical reality capable of supporting both points and spans.

Dimension: A representational framework used to locate, measure, compare, or classify selected aspects of reality.

The distinction developed here provides the foundation for the concepts introduced in the following sections. Variables assign values to dimensions, states are represented by collections of such values, and systems are subsequently analysed using these symbolic representations. Throughout this process, the symbolic representations should not be confused with the underlying reality that they describe.

2.4 Variables and Values

Properties may remain relatively constant or may differ between entities and change through time. To describe and compare such variation, properties may be represented by mapping them onto property dimensions. The concepts of spatio-temporal dimensions, property dimensions, variables, and values provide a framework for these representations.

Concrete entities exist within space and time. Consequently, they occupy positions and extents within the familiar three spatial dimensions and one temporal dimension, collectively referred to as spatio-temporal dimensions. Space and time are ontological aspects of reality. The dimensions used to represent them are epistemic constructs. Different observers may represent the same reality using Cartesian, polar, cylindrical, spherical, or other coordinate systems, with different orientations and origins, without altering the underlying reality being represented.

Spatio-temporal dimensions therefore provide epistemic representations of the locations and extents occupied by concrete entities within space and time. Unlike property dimensions, they permit entities to occupy spans rather than merely single positions. A tree, for example, occupies a range of positions within each of three spatial dimensions and persists over an interval of time. It does not occupy a single point within space-time.

Property dimensions perform a different representational role. Rather than representing where or when an entity exists, they represent variation with respect to selected properties. At a specified time, an entity is represented as occupying a single position on the property dimension associated with that property.

The distinction may therefore be expressed simply:

Spatio-temporal dimensions represent the locations and extents occupied by entities within space and time. Property dimensions represent the positions occupied by entities with respect to selected properties.

Some properties are derived from an entity's spatio-temporal extent or organisation. Consider the property of height. A tree occupies a vertical spatial span. The extent of that span may be recognised as the property height and represented by a single position on a height dimension. A complex spatio-temporal representation has thereby been compressed into a single property representation.

Other properties do not arise from spatio-temporal extent in this way. The mass of an entity, for example, reflects the combined organisation of its components rather than its spatial or temporal extent. Nevertheless, it too may be represented by a single position on a mass dimension. Similar representations apply to temperature, pressure, density, trust, resilience, flexibility, justice, and countless other properties.

Property dimensions therefore provide a means of comparing entities without representing their complete spatio-temporal organisation. Entities may be arranged according to increasing height, mass, temperature, age, trust, or other properties. The entities themselves remain discrete, while the property dimension may be continuous, allowing indefinitely many intermediate positions to be represented between any two positions.

The underlying variation represented by a property dimension is assumed to be real. Entities physically differ in height, mass, temperature, and many other respects. The recognition of a particular form of variation as a property, however, involves cognitive recognition and compression. A property dimension is therefore an epistemic representation of an ontological variation in reality.

To communicate and compare positions more precisely, positions on property dimensions may be mapped onto appropriate symbolic numerical systems. Different properties naturally employ different numerical domains. Quantities are commonly represented using the natural numbers or integers, continuous measurements using the real numbers, probabilities using the interval between zero and one, logical values using Boolean representations, and many other mathematical domains are employed where appropriate. Mathematics therefore comprises a family of symbolic representational systems rather than a single universal numerical language.

Within a chosen numerical domain, a measurement scale (Stevens, 1946) provides a conventional mapping between positions on the property dimension and symbolic numerical values. Height, for example, may be represented using metres or feet, mass using kilograms or pounds, and temperature using degrees Celsius or Fahrenheit. The underlying property does not depend upon the particular scale selected. The scale merely provides a conventional symbolic representation suitable for communication and quantitative reasoning.

The distinction between the underlying property and its symbolic representation is important. A tree does not become taller when its height is expressed in centimetres rather than metres, nor heavier when expressed in pounds rather than kilograms. Only the symbolic representation changes. The physical variation is ontological. The property dimension provides an epistemic representation of that variation. The numerical domain and measurement scale provide further symbolic representations that enable comparison, calculation, communication, and reasoning.

Not all property dimensions can be mapped easily or reliably onto numerical representations. We may recognise that one relationship involves greater trust than another, or that one organisation is more adaptable than another, without possessing a valid numerical scale capable of assigning precise values. Recognition of a property dimension therefore does not imply the existence of an appropriate numerical representation.

A variable is a property whose position on a property dimension may differ between entities or change through time. The value of a variable is the position attributed to a particular entity on that property dimension at a specified time. Where appropriate, that position may be represented numerically using an agreed mathematical domain and measurement scale.

Measurement uncertainty should not be confused with variation in the property itself. An observer may estimate the height of a tree as lying between 12.3 and 12.5 metres. This range represents uncertainty concerning the tree's position on the height dimension. It does not imply that the tree simultaneously possesses every height within that range.

The definitions adopted in this framework are therefore:

Property Dimension: An epistemic representation of the possible positions that entities may occupy with respect to a selected property.

Variable: A property of an entity whose position on a property dimension may differ between entities or change through time.

Value: The symbolic representation of the position attributed to an entity on a property dimension at a specified time.

The assignment of values to properties may therefore be understood as a sequence of epistemic compressions. Physical reality is first recognised in terms of entities and their properties. Selected properties are then represented by positions on appropriate property dimensions. Those positions may subsequently be represented within suitable mathematical domains and, where appropriate, mapped onto conventional measurement scales (Stevens, 1946) and symbolic numerical values. Each stage preserves those aspects of reality relevant to the analytical purpose of the investigation while discarding unnecessary detail. The resulting symbolic representations enable properties to be compared, communicated, analysed, and manipulated without requiring the underlying reality to be represented in its entirety.

Mathematics provides one of the most powerful families of symbolic representational systems developed for this purpose. Its effectiveness derives from its ability to preserve important structural and quantitative relationships present in reality while representing them in forms suitable for analysis, reasoning, prediction, and communication. Mathematical representations should therefore not be confused with the physical reality they describe. Rather, they provide highly compressed symbolic representations whose usefulness depends upon the extent to which they faithfully preserve the relevant relationships present within that reality.

The variables and values associated with an entity provide a means of describing its state. At a specified time, an entity may be represented by its positions across a selected set of property dimensions. Together these dimensions define a multidimensional representational space within which different possible states of the entity may be distinguished. The concept of state is considered next.

2.5 State

At any specified time, an entity possesses a particular combination of properties and values. This combination describes the state of the entity (von Bertalanffy, 1968).

Consider a tree. At a particular time, it has a particular height, mass, age, temperature, water content, and many other properties. Each variable may be represented as occupying a position on its associated property dimension. Taken together, these positions provide a multidimensional representation of the tree's state.

In practice, it is rarely possible or useful to represent every property of an entity. The properties selected depend upon the analytical purpose of the investigation. A forester, ecologist, and structural engineer may describe the state of the same tree using different variables. Each representation is therefore a compressed description of the entity rather than a complete representation of its physical reality.

This does not imply that the state of the entity is created by the observer. The entity possesses physical properties independently of their observation. What is epistemic is the selection of properties and values used to represent its state.

The definition adopted in this framework is:

State: The set of properties and associated values that apply to an entity at a specified time.

A state may therefore be represented as a position within a multidimensional space formed by the relevant property dimensions. If the value of one or more variables changes through time, the entity undergoes a change of state. This is considered next.

2.6 Change and Change of State

Entities exist through time. Where the value of one or more of an entity's variables differs between two specified times, the entity has undergone a change of state.

Consider a tree whose height increases from one year to the next. At each specified time, the tree occupies a position on the property dimension of height. A difference between those positions represents a change in the value of the variable and therefore a change in the state of the tree. The same principle applies where temperature rises, mass decreases, spatial position alters, or any other relevant property changes.

A change of state may therefore be represented as movement between positions within the multidimensional state space introduced in the preceding section. If the selected variables and their values remain unchanged, the represented state remains unchanged. If one or more values change, the entity occupies a different position in state space.

The identification of change necessarily involves comparison through time. A single observation may describe a state, but change can only be recognised by comparing states at different times. The time interval selected may affect the change that is observed. A tree may appear unchanged over several seconds while undergoing substantial change over several decades. Change is therefore always considered relative to a temporal interval and the variables selected for observation.

Changes of state may be gradual or abrupt. Where a variable changes progressively through intermediate values, the change may be described as continuous. In other cases, an entity may be represented as moving between distinguishable states, producing a discrete change of state. The distinction may partly reflect the nature of the underlying change and partly the scale, resolution, and variables used in its representation.

It is useful to distinguish change as the general occurrence of difference through time from change of state as its representation in terms of an entity's properties and values.

A change in physical reality may occur without being observed or represented. A change of state is identified when differences between the states of an entity at different times are recognised.

The definitions adopted in this framework are therefore:

Change: A difference in physical reality through time.

Change of State: A change in one or more of the properties or values that apply to an entity between specified times.

Continuous Change of State: A change of state in which a variable passes progressively through intermediate positions on a property dimension.

Discrete Change of State: A change of state in which an entity is represented as moving between distinguishable positions or states without intermediate positions being represented.

These concepts provide the temporal foundation for considering relationships, interactions, events, and processes. In particular, where change in one entity is associated with causal interaction with another, the concepts developed here allow that interaction to be represented in terms of changes of state (von Bertalanffy, 1968).

3. Relationships, Interactions and Processes

3.1 Relationships

Entities do not exist merely as isolated portions of reality. They occupy positions relative to other entities and may also interact with them. Systems theory therefore depends heavily upon the concept of a relationship.

In ordinary language, the term *relationship* is used very broadly. Two objects may be described as being close together, one person may be related to another, two variables may be statistically correlated, and one event may be said to cause another. These relationships are not necessarily of the same kind. It is therefore useful to distinguish between two broad forms of relationship: configurational relationships and causal relationships.

A configurational relationship arises from the arrangement of entities in space-time. One entity may be above, below, beside, within, or distant from another. Two events may occur simultaneously or at different times. Such relationships describe the relative configuration of entities but do not, by themselves, imply that matter, energy, or information passes between them.

For example, two stones lying beside one another possess a configurational relationship. Their relative positions can be described even if neither affects the other. Similarly, one star may appear closer to another in the night sky without any significant causal interaction occurring between them. The relationship arises from their relative disposition in space-time.

A causal relationship, by contrast, exists where entities affect one another through the transfer of matter, energy, or information. Within this series, information is understood as organised or structured matter or energy, and is therefore treated as a physical phenomenon capable of causal transfer. A flame heats a container through energy transfer; water entering soil is a transfer of matter; and a signal passing between two organisms transfers information. In each case, the relationship involves a physical transfer capable of contributing to a change of state.

The distinction is important because configurational relationships and causal relationships may appear similar in diagrammatic representation. A diagram may connect two entities with a line, while ordinary language may simply state that two things are "related". Yet a line indicating spatial proximity represents something fundamentally different from an arrow indicating causal transfer.

Configurational and causal relationships may also be associated. The configuration of entities may enable or inhibit causal interaction. Two entities may need to be sufficiently close before energy or matter can pass between them, while the structure of a network may determine which entities are capable of exchanging information. Configuration can therefore affect the possibility of causal relationships without itself constituting a causal transfer.

Relationships are treated in this framework as physically instantiated aspects of reality rather than purely conceptual associations. A spatial relationship is instantiated in the positions of entities within space-time. A causal relationship is instantiated through the physical processes by which matter, energy, or information is transferred. Our description of a relationship is epistemic, but the relationship described may exist independently of its observation.

The definitions adopted in this framework are:

Relationship: A physically instantiated circumstance in which one entity may be described with reference to another through either their relative configuration or the transfer of matter, energy, or information between them.

Configurational Relationship: A relationship arising from the configuration of entities in space-time, without transfer.

Causal Relationship: A physically instantiated process involving the transfer of matter, energy, or information between entities.

These distinctions provide the basis for considering interaction. Where entities enter into causal relationships and affect one another's states, they may be said to interact.

3.2 Interaction

Where entities participate in a causal relationship involving the transfer of matter, energy, or information, they may be said to interact. Interaction therefore involves more than entities merely being related or arranged with respect to one another. It involves causal transfer.

Consider two books resting side by side on a shelf. Their relative positions constitute a configurational relationship, but they need not interact causally. If one book is pushed so that it strikes the other, momentum and energy are transferred between them. The relationship has changed from being purely configurational to involving a causal interaction.

Causal interactions involve the transfer of matter, energy, or information. Fuel flowing into an engine transfers matter; sound travelling from a loudspeaker to a listener transfers energy; and a digital message sent between two computers transfers information. Such transfers may contribute to a change in the state of an entity, but a measurable change of state is not required for an interaction to occur. Continuous transfers may persist even when the receiving entity remains in a stable state.

Continuous causal transfers may instead maintain an existing state by balancing or counteracting other transfers. A building maintained at a constant temperature, for example, may remain in the same represented thermal state despite continuous flows of energy into and out of it. The absence of an observed change of state therefore does not imply the absence of causal interaction. A stable state may be the product of continuing causal activity.

The distinction between causal relationship and interaction requires some care. A causal relationship describes the physically instantiated process through which transfer occurs. Interaction describes the participation of entities in that causal relationship. The concepts therefore emphasise different aspects of the same causal organisation.

Interactions need not be reciprocal. Entity A may transfer matter, energy, or information to entity B without B producing a significant causal effect upon A. Where causal transfers occur in both directions, the interaction may be described as reciprocal. Many systems contain complex networks of both one-way and reciprocal interactions.

The consequences of an interaction may depend upon the state and organisation of the entities involved and the circumstances in which the interaction occurs. The same transfer may contribute to different changes in different entities, or in the same entity at different times. Its effects may also be balanced by other transfers. Interaction should therefore not be understood as a fixed relationship between a transfer and an inevitable observable change.

Interaction is fundamental to systems theory because systems are not merely collections of entities. Their organisation depends upon causal interactions between components. Recurrent or continuing interactions may generate and maintain persistent structures, processes, and patterns of change. These ideas are developed later in this paper.

The definition adopted in this framework is:

Interaction: Participation by two or more entities in a causal relationship involving the transfer of matter, energy, or information between them (von Bertalanffy, 1968).

Interaction therefore links the concepts developed so far. *Entities possess properties and states; causal relationships involve transfer; and interaction occurs where entities participate in those relationships. Such interaction may change an entity's state or contribute to its maintenance.*

The next concept, event, considers how causal interactions are represented within bounded periods of time. While some interactions are naturally brief, others continue over extended periods. Events provide a means of identifying and representing temporally bounded causal activity, including bounded intervals within ongoing interactions.

3.3 Events

Causal relationships occur through time. Some interactions are brief and naturally bounded, while others continue for extended periods. Where a causal interaction, or a distinguishable part of an ongoing causal interaction, is identified within a bounded period of time, it may be represented as an event.

In ordinary language, events are usually understood as things that happen. A collision, an explosion, a conversation, or the arrival of a signal may each be described as an event. Such descriptions identify a causal occurrence and distinguish it temporally from the circumstances preceding and following it.

Events are therefore closely associated with change of state. A collision may alter the position or velocity of an entity, while the arrival of a signal may change the state of a receiving system. Many events are recognised precisely because an observable change occurs.

However, change of state should not be treated as a necessary condition for an event. As discussed in the preceding section, causal interactions may maintain an existing state through continuing transfers. An electric motor operating at a constant speed continuously receives electrical energy and transfers mechanical energy, even though its represented operating state remains unchanged. A period of steady operation may still be identified as an event, despite no change occurring in the selected state variables.

The identification of an event therefore involves temporal bounding. A causal interaction may continue for a long period, but an observer may identify a distinguishable interval within that ongoing interaction for the purposes of description or analysis. A storm may persist for several days, while a particular lightning strike is identified as a shorter event within it. Similarly, a continuous flow of water may be represented as a series of operational events when a valve is opened, adjusted, or closed.

The boundaries of an event may sometimes be physically conspicuous. A collision begins when entities make causal contact and ends when the particular interaction ceases. In other cases, temporal boundaries are selected for analytical purposes. A conversation may be treated as one event, or individual exchanges within it may be represented as separate events.

Events therefore exhibit a form of set duality similar to that considered earlier. An event may be treated as a single entity or disaggregated into a collection of component events and causal interactions. The level at which an event is identified depends partly upon the organisation of the causal activity and partly upon the purpose and perspective of the observer.

This does not imply that events are merely subjective constructions. The causal interactions and transfers represented by an event are physically instantiated in space-time. What may vary is the temporal and analytical boundary used to identify them as a particular event.

The definition adopted in this framework is:

Event: A temporally bounded causal interaction, or organised set of causal interactions, treated as an identifiable entity.

Events represent one way of describing causal activity through time. Where attention shifts from temporal boundaries to the organisation, continuity, or transformation of causal activity, the same underlying reality may instead be represented as a process. The concept of process is considered next.

3.4 Processes and Systems

Events identify causal activity within bounded periods of time. Processes, by contrast, emphasise the continuity and organisation of causal activity through time. Where ongoing causal interactions are considered in terms of the transformations they produce, sustain, or regulate, they may be represented as processes (von Bertalanffy, 1968).

A process receives one or more inputs, acts upon or transforms them, and produces one or more outputs (Hall & Fagen, 1956). A plant may receive water, carbon dioxide, and solar energy and produce growth, oxygen, and other outputs. A manufacturing process receives materials and energy and produces manufactured goods and waste. A cognitive process may receive information and produce a decision or other informational output.

A process may therefore be represented in simplified form as:

Inputs → Process → Outputs

This representation emphasises transformation. The process is identified not simply because causal interaction occurs, but because an organised sequence of causal interactions transforms incoming matter, energy, or information into outgoing matter, energy, or information.

The concept of process is closely related to that of function. When a process is considered in terms of what it does, its function may be described by the transformation it performs. For example: a pump moves fluid; a filter separates selected material from a flow; and a translation process transforms information from one symbolic representation into another. Function and purpose are considered more fully later in this paper, but the relationship is worth noting here.

Processes may be continuous, intermittent, or represented as discrete sequences of events. A river continuously transports matter and energy; a pump may operate intermittently; and a manufacturing process may be represented as a sequence of distinct production cycles. These differences do not alter the basic concept. In each case, a process is identified through an organised causal transformation involving inputs and outputs.

Processes may be represented at different levels of detail. Photosynthesis may be treated as a single process for one purpose but disaggregated into multiple interacting biochemical processes for another. Similarly, manufacturing may be represented as one process or as an organisation of machining, assembly, inspection, and transport processes. As with events, the causal activity is physically instantiated, while the boundary and level at which a process is identified depend partly upon the purpose and perspective of the observer.

A process may be represented as a single organised causal transformation. However, processes may also interact. The outputs of one process may become inputs to another, creating an organised network of causal transformations. Such a network does not necessarily constitute a system. Emergent organisation is recognised through the novel causal outputs that it enables. If an organised network of interacting processes is incapable of producing causal outputs beyond those of its component processes acting independently, there is no operational basis for recognising it as a distinct system. A system therefore exists where the organisation of interacting processes gives rise to emergent outputs, or the capability to produce such outputs, that cannot be produced by the component processes acting independently. These emergent outputs may include new behaviours, functions, products, services, or other causal effects. The organised network may then be represented as a system.

A simple system may comprise a single organised process. A pump, for example, may be represented as a system receiving matter and energy as inputs and producing a flow of matter and other outputs. At a finer level of analysis, however, the same pump may be represented as several interacting mechanical and physical processes. Whether it is represented as a single process or as an organisation of processes therefore depends partly upon the purpose and perspective of the observer.

More complex systems may comprise many interacting processes. The outputs of one process may provide inputs to another, while other inputs may be received from outside the system. The system is distinguished not merely by the presence of multiple processes but by their causal organisation. Their interactions form an identifiable whole that may exhibit properties, behaviours, or outputs associated with the organisation rather than with any process considered in isolation.

For scientific investigation, systems are represented as wholes according to the analytical purpose of the observer. The phrase represented as a whole is important. The causal organisation is assumed to exist in physical reality, but its recognition and representation as a system are epistemic. The same causal organisation may be represented at different levels or with different boundaries according to the purpose of investigation. A process treated as a single system at one level may be disaggregated into several interacting processes at another.

The distinction between a process and a system is therefore primarily one of emphasis, but they are not equivalent concepts. A process foregrounds causal transformation: what enters, what occurs, and what leaves. A system foregrounds causal organisation: the interacting processes involved, the relationships between them, and the emergent properties, behaviours, or outputs arising from that organisation. The same portion of reality may sometimes be represented in either way, but the representations draw attention to different aspects of its causal structure.

The relationships between processes within a system may differ in the extent to which they depend upon one another. A process whose inputs are derived primarily from the outputs of other processes within the same system is closely coupled to the internal organisation of that system. A process whose inputs are derived primarily from outside the system is more loosely coupled to its internal processes. Many processes depend upon a combination of internal and external inputs and therefore occupy intermediate positions between these conditions.

Coupling should therefore not be regarded simply as the presence or absence of interaction. It concerns the degree to which a process depends upon causal inputs generated by other processes within the same system rather than upon inputs derived from outside it. This distinction becomes particularly important when considering system environments, hierarchical organisation, and emergence later in this paper.

The definition adopted in this framework are therefore:

Process: An organised causal transformation in which inputs of matter, energy, or information are acted upon to produce outputs of matter, energy, or information.

When an organised assembly develops identifiable emergent properties, behaviours, or outputs, it is recognised as a system. This distinction is developed later in this paper. The concepts of system boundaries, systems of interest, and representation are also developed later in this paper.

3.5 Transfers

Causal relationships, interactions, and processes depend upon transfer. Matter and energy may pass between entities, while information may be transmitted through structure embodied in matter or energy. Within the framework developed here, transfer is therefore regarded as a physical phenomenon (Shannon, 1948).

For explanatory convenience, this paper refers to transfers of matter, energy, or information. Information is not, however, treated as a separate non-physical substance. Informational transfer involves matter or energy whose organisation is capable of contributing to causal effects in a receiving entity.

That which is transferred may be simple or highly organised. A transfer may comprise matter, energy, a structured signal, or an entity that is itself a causal structure, process, or system. In some circumstances, the whole of an identified system may be transferred into a different environment while retaining its internal causal organisation. It may then interact with other entities and systems and contribute to the formation of a larger causal organisation or assembly.

The consequences of a transfer depend not only upon what is transferred but also upon the organisation and state of the recipient. A radio signal may physically reach both a receiver and a stone, but only the receiver possesses the causal organisation required to respond to the signal as information. The same transfer may therefore produce different effects in different recipients.

Transfers may be continuous, intermittent, or discrete. They may occur between processes within a system or across a selected system boundary. Whether a transfer is described as internal, an input, or an output depends upon the boundary of the system being represented. The underlying causal transfer remains the same; its systems description changes with analytical perspective.

The definition adopted in this framework is:

Transfer: The physical passage or transmission of matter, energy, or information through a causal relationship.

The nature of causal transfer, the relationship between source and recipient, the limits upon what may be transferred, and the causal amplification of small transfers are examined more fully in a subsequent paper on Foundational Causality.

4. Configuration, Organisation and Structure

4.1 Introduction and Terminological Convention

Throughout this paper a distinction is maintained between physical reality, its cognitive recognition, and its symbolic representation. Many systems concepts describe ontological phenomena, while others describe the ways in which those phenomena are recognised and represented. Confusing these levels has contributed to ambiguity in systems terminology.

Natural language frequently uses closely related words to describe processes, properties, and entities. In systems theory this can also lead to ambiguity because the same word may refer to an activity, an abstract concept, or an identifiable exemplar. For clarity, the following conventions are adopted throughout this paper.

Finally, Table 1 presents the principal terminology used in this paper to represent progressively organised ontological arrangements. The sequence reflects a hierarchy of increasing organisation in which each successive concept subsumes those preceding it while introducing additional recognisable characteristics. A set is formed by aggregating entities or components. When these possess a recognisable arrangement in space and time they form a configuration. When configured entities become integrated into a coherent whole they form an assembly. Finally, when the components or subsystems of an assembly interact in a coordinated causal manner that gives rise to identifiable emergent properties, behaviours, or outputs, the assembly is recognised as a system. Thus, every system is necessarily an assembly, every assembly possesses a configuration, and every configuration is an aggregate, although the converse is not necessarily true.

The terminology presented in the table should be understood as an epistemic framework for representing progressively organised reality rather than as reality itself. The concepts are models used to describe recurring patterns of organisation observed in the world. As organisation increases, increasingly rich concepts become available to represent it. At the same time, higher-order concepts remain compatible with those below them because they incorporate their characteristics while adding further distinguishing features. This hierarchical relationship between organisation and terminology provides a coherent foundation for the Systems Linguistics developed in other papers where it is shown that such terms represent progressively richer conceptual compressions of the underlying ontological arrangements.

	Set[2]	**Configuration**[2]	**Assembly**[2]	**System**[2]
	Recognisable features	+ recognisable arrangement in space-time	+ recognisable causal interactions	+ recognisable emergent properties
Noun	An entity	An entity	An entity	An entity
Noun (relative to collection)	A component	A component	A component	A component or sub-system
Collective Noun	A collection of entities or components	A collection of entities or components	A collection of entities or components	A collection of components or sub-systems
Verb	To aggregate a collection of entities or components	To configure a collection of entities or components	To assemble a collection of entities or components	To organise a collection of components or sub-systems
Conceptual Term	A set or aggregate	A configuration	An assembly[1]	A system
Recognising Characteristic	Aggregation[3]	Configuration[3]	Assembly[3,4] or structure[3,4]	Organisation[3,4] or structure[3,4]

Notes

1) Assemblies may be pre-systemic.
2) Each higher-order organisational concept incorporates the characteristics of those below it while adding new distinguishing characteristics.
3) Aggregation, configuration, assembly, organisation and structure may be static or dynamic.
4) The terms *assembly* and organisation emphasise the relationships between components whilst *structure* emphasises their overall pattern within the whole.
5) The terms *structure* and *organisation* may also be used generically to describe aggregation, configuration, assembly, or organisation.

Table 1. Terminology used to represent progressively organised ontological arrangements.

4.2 Collections and Sets

The preceding sections considered identifiable entities and the relationships and causal interactions between them. Systems theory, however, frequently requires several entities to be considered together. This introduces the concepts of collections and sets.

A collection is a plurality of entities considered together on the basis of some shared feature, circumstance, or condition of selection. The entities need not interact causally or form an organised whole. The basis for considering them together may be no more than spatial proximity, temporal coincidence, or inclusion within a selected spatio-temporal boundary. A group of stones lying within an area of beach, a list of people, a group of stars, or a selection of events may each be treated as a collection.

The same entities may also be represented as a set. The distinction lies primarily in representational perspective. When the constituent entities are foregrounded and considered as a plurality, they are represented as a collection. When the plurality is treated as a single entity, it is represented as a set.

A group of birds occupying the same region of sky, for example, may be considered as a collection of individual birds. The same birds may also be treated as a single flock. The physical entities and the basis on which they are considered together remain unchanged. What changes is whether plurality or unity is foregrounded.

This distinction provides an example of complementary representations. To ask whether selected entities are a collection or a set may impose an unnecessary choice. The same portion of reality may support both representations. The concept of a collection foregrounds the constituent entities; the concept of a set foregrounds them as a single entity.

Collections and sets may be identified on the basis of properties, variables, values, relationships, or other circumstances. A group of trees, for example, may be identified because its members possess the distinguishing properties associated with trees. A group of objects exceeding one metre in height may be identified because their positions on the relevant property dimension satisfy the selected condition. Similarly, a group of events may be identified because they occurred within a selected period or location.

The identification of collections and sets is therefore closely related to the concepts of properties, variables, and values introduced earlier. A property may provide the basis for grouping entities, while a value or range of values on a property dimension may provide a condition for inclusion. The condition "height greater than one metre", for example, allows the relevant entities to be identified. Once identified, the same group may be represented either as a collection, emphasising the individual members, or as a set, emphasising the group as a unified whole.

The dual representation of a selected plurality as either a collection or a set gives rise to set duality. A set may be represented as a single entity while simultaneously being understood as a plurality of constituent entities. The set *trees*, for example, may be treated as the singular entity *trees* or disaggregated into the individual trees. This duality is important for systems theory.

However, a set should not automatically be regarded as a system. The members of a set may share properties or circumstances without interacting causally. The set of all red cars, for example, is identified through a shared property, but the individual cars do not thereby form a single causal organisation. Similarly, a collection of machine components placed in a box does not become a machine merely because the components are considered together or represented as a set.

This distinction returns us to the difference between aggregational hierarchy and organisational hierarchy. Collections and sets support aggregation: entities are considered together because they share selected properties or circumstances. Systems require causal organisation: one or more processes and their interactions are represented as an organised whole. A set may contain systems, and the components of a system may be represented as a set, but set membership alone does not establish systemhood.

The definitions adopted in this framework are therefore:

Collection: A plurality of entities considered together on the basis of a shared feature, circumstance, or condition of selection.

Set: A collection represented as a single entity.

Set Duality: The principle that the same plurality of entities may be represented either as a collection of constituent entities or as a single set, depending upon analytical perspective.

Set Duality provides a connection between the ontology developed here, set theory and formal logic. Groups of entities may also be represented formally. Set theory represents groups as sets and provides a formal means of describing relationships between them (Cantor, 1895). Formal logic primarily treats groups as collections of entities and provides a means of expressing and reasoning about them. The two approaches are complementary and widely used throughout science and mathematics. Systems theory asks the additional question of whether and how the entities concerned are causally organised.

That final question leads naturally to the concept of configuration.

4.3 Configuration

A collection or set identifies which entities are being considered. A configuration describes how those entities are arranged with respect to one another.

Configuration refers to the spatial, temporal, or spatio-temporal arrangement of entities and the relationships arising from that arrangement. A row of houses, the arrangement of planets within a solar system, the pattern of branches within a tree, or the relative positions of components within a machine are all examples of configurations.

Configuration does not, in itself, imply causal interaction or organisation. The entities forming a configuration may or may not interact. A group of parked cars occupying particular parking spaces possesses a configuration because the cars occupy particular positions relative to one another, even if no significant causal interaction occurs between them.

Configuration should therefore be distinguished from collection. A collection identifies which entities are being considered. A configuration identifies how those entities are arranged.

Configurations may be static or dynamic. A static configuration describes the arrangement of entities at a particular time. A dynamic configuration describes how that arrangement changes through time. A flock of birds in flight, traffic moving through a road network, or the changing positions of planets may each be represented as dynamic configurations.

Configuration is closely related to the configurational relationships introduced earlier. Those relationships describe the relative positions of entities. Configuration represents the collective pattern formed by those relationships. Individual configurational relationships therefore combine to produce an overall configuration.

Configurations may also be represented at different levels of detail. A city may be represented simply as a collection of districts, more fully as a configuration of roads and buildings, or in much finer detail as individual structures and their internal arrangements. The physical reality remains the same. What changes is the level of representation selected by the observer.

Configuration is an important concept because recurring configurations often provide the basis for recognition. As discussed in Section 2.1, concrete entities are commonly recognised through recurring configurational similarities that may be cognitively compressed into prototypes. Configuration therefore links the physical arrangement of reality with the cognitive processes by which entities are recognised and represented.

Configurations may contribute to, enable, or constrain causal interaction. The relative positions of entities may determine whether matter, energy, or information can be transferred between them and may therefore influence the causal organisation that subsequently develops. Configuration is consequently a necessary foundation for organisation but should not itself be confused with organisation.

The definition adopted in this framework is therefore:

Configuration: The overall spatial, temporal, or spatio-temporal arrangement of a plurality of entities.

A configuration may be understood as the collective pattern formed by the configurational relationships between entities. Individual relationships describe how one entity is situated with respect to another. Configuration describes the overall arrangement that emerges from those relationships.

Configuration provides the foundation for organisation. Where recurring causal interactions develop within or between configurations, organised structures and systems may emerge.

4.4 Assemblies

The preceding sections introduced collections, sets, and configurations. An assembly builds upon these concepts by introducing causal interaction.

A collection is a plurality of entities considered together. A configuration describes the arrangement of those entities in space and time. An assembly is formed where the entities participate in causal interactions and may therefore be regarded as an identifiable whole.

Assemblies may exist at many levels of complexity. An interacting collection of machine components forms an assembly. A biological cell is an assembly of interacting molecules and organelles. A tree is an assembly of interacting cells and tissues. Similarly, a human society may be represented as an assembly of interacting individuals and organisations. The concept therefore applies equally to physical, biological, and social systems.

An assembly should not be confused with a set. A set represents a plurality of entities as a single abstract entity for the purposes of classification, comparison, or reasoning. An assembly is an identifiable portion of physical reality whose constituent entities participate in causal interactions. Every assembly may be represented as a set, but not every set represents an assembly.

Assemblies may also be represented at different levels of analysis. A tree may be represented as an assembly of cells, as a component within a forest ecosystem, or as part of a wider ecological assembly. Similarly, a manufacturing plant may be represented as an assembly of production processes, while itself forming part of a larger industrial assembly. The level at which an assembly is recognised depends upon the purpose and perspective of the observer.

The concept of assembly provides a convenient general term for an identifiable interacting whole without implying any particular level of complexity, persistence, or emergence. It therefore avoids some of the ambiguity associated with the term organisation, which in ordinary English commonly refers to an interacting group of people. Throughout this paper, the broader term assembly is therefore used where a general systems concept is intended.

The definition adopted in this framework is therefore:

Assembly: An identifiable entity comprising a plurality of causally interacting constituent entities.

Assemblies possess many properties, including configuration, organisation, coupling, persistence, complexity, and, where sufficient organisation develops, emergent properties. The following section considers organisation as one of the fundamental properties of assemblies.

4.5 Pre-Systemic Assemblies and Systems

The concepts developed so far describe progressively richer forms of organisation. Collections identify pluralities of entities, configurations describe their arrangement, assemblies identify causally organised entities, and systems identify organised assemblies exhibiting emergent properties, behaviours, or outputs. Organisation refers to the ontological pattern of causal interactions, while structure refers to the epistemic representation of that organisation.

Not all structured assemblies exhibit emergent properties, behaviours, or outputs. Many assemblies possess identifiable organisation and structure while behaving largely in the same manner as their constituent entities. Such assemblies are referred to here as pre-systemic assemblies.

A pre-systemic assembly is therefore an assembly possessing organisation and structure, but whose organisation is insufficient to generate identifiable emergent properties, behaviours, or outputs. Although causal interactions occur between its constituent entities, those interactions do not yet give rise to new causal capabilities attributable to the assembly as a whole.

The distinction is one of degree rather than kind. As the organisation of an assembly develops, new interactions become possible, existing interactions become increasingly integrated, and constraints upon the behaviour of constituent entities become more extensive. At some point, these developments may permit the emergence of properties, behaviours, or outputs that cannot reasonably be attributed to the constituent entities considered independently. The assembly may then be recognised as a system.

This transition should not be regarded as a sharp boundary in physical reality. The development of organisation is typically gradual, and different observers may identify the emergence of systems at slightly different stages according to the phenomena under investigation. Nevertheless, the appearance of identifiable emergent properties provides an objective basis for distinguishing systems from pre-systemic assemblies.

This interpretation also clarifies the relationship between assemblies and constraints. Every assembly possesses configuration, organisation, and structure, each of which constrains the possible interactions of its constituent entities. As assemblies become more highly organised, these constraints become increasingly integrated, enabling new causal interactions and eventually the emergence of new properties, behaviours, or outputs. Systems may therefore be understood as assemblies whose organisation of constraints has become sufficient to generate emergence.

The definitions adopted in this framework are therefore:

Pre-Systemic Assembly: An assembly possessing organisation and structure but not exhibiting identifiable emergent properties, behaviours, or outputs.

System: An organised assembly exhibiting identifiable emergent properties, behaviours, or outputs and recognised and represented as a whole.

The distinction between a pre-systemic assembly and a system is therefore explained not by the addition of a new kind of entity, but by the emergence of identifiable properties, behaviours, or outputs arising from causal organisation. Pre-systemic assemblies often provide the organised foundations from which higher-level systems emerge through the further development and integration of causal organisation. The nature of emergent properties and the processes by which they arise are examined in the next section.

4.6 Progressive Organisation and Representation

The concepts introduced in this section do not represent independent categories but progressively richer forms of organised reality. Every system is necessarily an assembly, every assembly possesses a configuration, and every configuration is an aggregation of entities. Each successive concept therefore subsumes those that precede it while introducing additional characteristics that justify its recognition as a distinct concept.

The progression from aggregation to configuration, assembly and system reflects increasing ontological organisation. Aggregation recognises only that entities are considered together. Configuration adds recognisable spatial or temporal arrangement. Assembly introduces recurring causal interactions between the entities. System adds the emergence of properties, behaviours, or outputs that are not exhibited by the individual components acting independently. Each stage therefore represents an increase in the richness of causal organisation.

This hierarchy is equally reflected in the corresponding processes. To configure entities first requires that they be aggregated. To assemble them requires that they possess a configuration. To organise them into a system requires that they first form an assembly. Thus the processes of aggregating, configuring, assembling and organising are themselves hierarchically related, with each higher process incorporating those beneath it.

These concepts are frequently used differently across disciplines and in everyday language. In ordinary usage, the terms *organisation* and *structure* are often employed as convenient generic descriptions for any recognisable arrangement of entities. A collection of parked vehicles, for example, might be described as being "well organised" or having a "clear structure" even though it exhibits neither significant causal interaction nor emergent behaviour. Such usage is entirely appropriate in informal communication but can lead to ambiguity in systems theory.

Within this framework, a distinction is therefore made between the generic and technical use of these terms. Organisation, in its technical sense, refers specifically to the recurring causal relationships that enable an assembly to function as a coherent whole. Structure refers to the recurring pattern or arrangement through which that organisation is realised (Simon, 1962). Both concepts therefore presuppose a degree of causal organisation beyond that of simple configurations or unorganised assemblies, although in generic language they may legitimately be used to describe any progressively organised collection.

The hierarchy presented in Table 1 is therefore both ontological and epistemological. Ontologically, it represents progressively richer forms of organisation within reality. Epistemologically, it represents progressively richer concepts developed to recognise and describe those forms. The terminology employed by different disciplines can consequently be understood as alternative compressions of this common underlying hierarchy, a point developed further in papers on Systems Linguistics.

The definitions adopted in this framework are therefore:

Organisation (Generic): The network of relationships or interactions that brings entities together into a recognisable whole. This applies equally to aggregates, configurations, assemblies and systems.

Structure (Generic): The overall pattern or arrangement exhibited by a network of relationships or interactions. Again, this applies throughout the hierarchy.

Organisation (Technical): The network of causal interactions between the constituent components or processes of an assembly or system.

Structure (Technical): The overall pattern exhibited by the causal organisation of an assembly or system.

4.7 Static and Dynamic Organisation or Structure

The forms of organisation described above may each be either **static** or **dynamic**. An aggregate, configuration, assembly or system may remain relatively unchanged throughout the period of observation or may undergo continual change while retaining its identity.

Similarly, the corresponding representations of these forms of organisation may exhibit static or dynamic structure. A static structure represents an organisation that changes little over the period of observation. A dynamic structure represents an organisation whose constituent entities, relationships or interactions change continuously while maintaining an identifiable pattern through time.

The distinction between static and dynamic applies throughout the hierarchy of organised collections rather than only to systems. What differs between the levels is not whether change occurs, but the richness of the underlying organisation. Aggregates exhibit only the coexistence of entities. Configurations additionally exhibit recognisable arrangement. Assemblies exhibit recurring causal interaction, while systems exhibit emergent properties, behaviours, or outputs arising from that organisation. At each level, both the organisation itself and its representation may be either static or dynamic.

A pile of bricks may be regarded as a largely static aggregate because the individual bricks are simply considered together and undergo little change during the period of observation. A row of parked cars exhibits a largely static configuration, since the vehicles occupy a recognisable spatial arrangement without necessarily interacting causally. A bridge represents a largely static assembly, as its components interact causally to maintain the integrity of the whole while its overall organisation remains relatively unchanged.

Other forms of organisation are inherently dynamic. The hands of a clock continually change position while maintaining a stable pattern of organisation that enables the clock to perform its function. Blood circulation within the human body is a dynamic system whose constituent processes continually interact to sustain the organism. Similarly, a transport network exhibits persistent patterns of organisation despite the continual movement of vehicles and passengers, while biological and social systems typically maintain their identities through ongoing patterns of interaction rather than through physical permanence. In each case, the underlying organisation persists even though the entities, relationships, or processes continually change.

The definitions adopted in this framework are therefore:

Static Organisation: Organisation that changes little over the period of observation.

Dynamic Organisation: Organisation whose constituent entities, relationships, or interactions change through time while maintaining an identifiable overall organisation.

Static Structure: Structure that changes little over the period of observation.

Dynamic Structure: Structure that changes through time while maintaining an identifiable pattern.

5. Recursive Organisation

5.1 The Recursive Nature of Systems

Throughout this paper, progressively richer forms of organisation have been developed from a small number of foundational concepts. Beginning with identifiable entities, the discussion introduced properties and collections, configurations, causal relationships or interactions and assemblies, and finally processes and systems exhibiting emergent properties, behaviours, or outputs. The concepts developed are not confined to a single level reality. Rather, they apply recursively.

Once an assembly is recognised as a system, that system may itself be treated as an identifiable entity. It may therefore participate in relationships and interactions with other entities, form part of collections and assemblies, and contribute to the emergence of larger systems. The same conceptual framework may consequently be applied repeatedly at successively higher or lower levels of organisation.

This recursive interpretation follows naturally from the distinction maintained throughout this paper between ontology and epistemology. Ontologically, systems are organised assemblies exhibiting emergent properties, behaviours, or outputs. Epistemically, once such a system has been recognised and represented, it may itself be treated as a single entity for the purposes of further investigation. The ontology has not changed; only the level at which reality is represented.

Reality may therefore be understood as exhibiting recursive organisation. Organised entities combine to form organised assemblies; organised assemblies exhibiting emergence are recognised as systems; and those systems may themselves become components of larger organised assemblies. The same progression from entity to system may therefore recur indefinitely across successive levels of organisation (Koestler, 1967; Simon, 1962).

This recursive character is one of the principal reasons why General Systems Theory is *general*. The concepts developed in this paper are intended to apply equally to physical, biological, cognitive, social, ecological, and technological systems. Although the entities, processes, and emergent properties differ between these domains, the underlying conceptual framework remains the same.

The following sections consider some of the consequences of this recursive organisation, including systems represented as components of larger systems, hierarchical organisation, and the choice of analytical level appropriate to a particular investigation.

5.2 Systems as Entities

A consequence of recursive organisation is that a recognised system may itself be treated as an entity. This does not imply that the internal organisation of the system has disappeared. Rather, it reflects the epistemic decision to represent the system as a coherent whole while retaining the possibility of examining its internal organisation whenever required.

Representing systems as entities allows progressively larger organisations to be understood without requiring every level of organisation to be considered simultaneously. For example, a biological cell may be represented as an entity when investigating the behaviour of tissues (Koestler, 1967). Similarly, an individual person may be represented as an entity when investigating social organisations, while a social organisation may itself be represented as an entity when investigating societies. At each level, the lower-level organisation remains real, but it is compressed into a higher-level representation appropriate to the analytical purpose of the investigation.

This recursive compression is essential for understanding complex systems. Without it, the number of entities and interactions requiring simultaneous consideration would rapidly become unmanageable. Treating systems as entities therefore provides a practical means of representing complexity while preserving the possibility of decomposing the representation whenever a more detailed investigation is required.

The process is entirely consistent with the systems linguistic principles introduced earlier in this work. A recognised system functions as a compressed representation of its internal organisation, just as words, symbols, diagrams, and models compress more detailed underlying structures. Productive investigation therefore involves moving fluently between compressed representations and the more detailed organisations they represent.

The same system may consequently be represented either as a complete system in one investigation or as a single entity within a larger system in another. The underlying reality has not changed; only the level of representation adopted by the investigator.

5.3 Hierarchical Organisation

The recursive formation of systems naturally gives rise to hierarchical organisation. As systems emerge and are recognised as entities, they may become components of larger organised assemblies, which in turn may exhibit their own emergent properties, behaviours, or outputs. This process may be repeated indefinitely, producing successive levels of organisation throughout natural and social reality.

Hierarchical organisation should not be understood as a rigid or purely classificatory structure. Rather, it reflects the repeated emergence of new causal organisations. Each higher level possesses properties, behaviours, or outputs that arise from the interactions of its components and cannot be fully explained by considering those components in isolation. At the same time, the continued existence and behaviour of higher-level systems remain dependent upon the organisation and interactions occurring at lower levels.

The resulting hierarchy is therefore one of causal organisation rather than simply one of containment. A biological cell may form part of a tissue, tissues may form organs,

organs may form organisms, and organisms may participate in populations, ecosystems, or societies. At each level, new forms of causal interaction and emergence become possible while remaining grounded in the organisation of the constituent systems.

Hierarchical organisation also reflects the practical requirements of scientific investigation. Different disciplines typically focus on different levels of organisation according to the phenomena they seek to explain. Physics investigates elementary physical systems, biology investigates living systems, psychology investigates cognitive systems, sociology investigates social systems, and ecology investigates environmental systems. Although the subject matter differs, each discipline investigates organised entities interacting within larger organised contexts.

General Systems Theory provides a common conceptual framework for understanding these different levels of organisation. Rather than treating each discipline as requiring fundamentally different systems concepts, the recursive perspective recognises that the same foundational principles apply throughout the hierarchy. Differences arise primarily from the nature of the entities, interactions, and emergent phenomena being investigated rather than from changes in the underlying systems concepts themselves.

5.4 Levels of Investigation

The recursive nature of systems provides investigators with considerable flexibility when selecting the level of organisation appropriate to a particular study. The same phenomenon may be investigated at different levels according to the questions being addressed, with each level providing insights that may not be apparent at others.

For example, the functioning of the human body may be investigated in terms of molecular interactions, cellular processes, organs, physiological systems, or the behaviour of the individual as a whole. Similarly, a business organisation may be investigated in terms of individual employees, departments, management structures, organisational culture, or its interactions with customers, suppliers, regulators, and wider society. Each perspective represents a valid level of investigation directed towards a different aspect of the same underlying reality.

The choice of level is therefore an epistemic decision rather than an ontological one. It reflects the analytical purpose of the investigation and the nature of the questions being asked, not a change in the reality being investigated. Higher-level investigations often compress the organisation of lower-level systems into identifiable entities, while lower-level investigations may decompress those entities to examine their internal organisation in greater detail.

No single level of investigation is inherently more fundamental than another. Understanding complex systems frequently requires movement between levels of organisation, relating higher-level emergent behaviours to the organisation of their constituent systems while recognising that higher-level systems may also influence the behaviour of their components through changing constraints and patterns of interaction.

The recursive framework developed throughout this paper therefore supports investigations across every level of organised reality. Whether attention is directed towards components, complete systems, or systems of systems, the same foundational concepts remain applicable. General Systems Theory consequently provides a common conceptual language capable of relating investigations conducted at different levels of organisation while respecting the distinct contributions made by individual disciplines.

5.5 Section Conclusion

This chapter has shown that the concepts developed throughout this paper apply recursively across successive levels of organised reality. Once recognised as coherent wholes, systems may themselves be represented as entities, allowing progressively larger organisations to emerge and be investigated using the same foundational concepts. The resulting hierarchy is therefore not one of changing systems principles, but of repeated application of those principles to increasingly complex forms of causal organisation.

The recursive perspective demonstrates why General Systems Theory is genuinely *general*. Whether investigating physical, biological, cognitive, social, technological, or ecological phenomena, the same concepts of entities, properties, relationships, interactions, organisation, emergence, systems, boundaries, environments, and behaviour remain applicable. What changes between domains is not the underlying conceptual framework but the nature of the entities, interactions, and emergent phenomena being investigated.

The distinction maintained throughout this paper between ontology and epistemology provides a coherent foundation for this recursive interpretation. Organised systems exist independently of their representation, while investigators may choose the level of organisation most appropriate to the questions being addressed. Higher-level representations compress the organisation of lower-level systems, yet those representations may always be decompressed when more detailed investigation is required.

General Systems Theory may therefore be understood as a unified conceptual framework for investigating organised reality. By providing a common language applicable across successive levels of organisation, it enables knowledge developed within different disciplines to be related, compared, and productively coordinated without requiring those disciplines to abandon their own specialised perspectives.

Subsequent papers will build upon these foundations. They will examine the causal interactions through which organised systems persist and change, and will consider how new systems emerge through the formation of novel causal organisations. Together they will extend the foundational concepts established in this paper towards a more comprehensive General Systems Theory.

6. System Properties

6.1 Introduction

The preceding sections developed the concepts required to understand systems from first principles. Beginning with entities, properties, relationships, and processes, progressively richer forms of organisation were introduced through collections, configurations, assemblies, organisation, structure, and the distinction between pre-systemic assemblies and systems. These concepts established the ontological foundations of systems together with the epistemic concepts used to recognise and represent them. Having established what constitutes a system, the present section examines systems themselves, including the selection of the system of interest, their components, their properties, and the role of emergence.

6.2 Representing Systems for Investigation

The underlying organisation of reality exists independently of any particular investigation. However, different investigations require different representations of that reality according to their explanatory purpose. Consequently, the same entity may be represented in different ways without implying any change in the underlying causal organisation.

The first step in systems investigation is to identify the system of interest. Having done so, the investigator must decide which surrounding entities should be represented as systems in their own right, which should be represented as subsystems, and which may be represented simply as components. This choice depends upon whether the emergent properties, behaviours or outputs of those entities are relevant to explaining the system of interest.

For example, when investigating the grasping capability of a human hand, bones, muscles, tendons and ligaments may be represented as components because their own biological organisation is not directly relevant to the investigation. The investigation focuses upon how their interactions collectively generate the emergent capability of grasping. In contrast, when investigating bone growth or fracture repair, a bone itself becomes the system of interest. Its cells, tissues and internal organisation are then represented explicitly because their emergent properties are now directly relevant to the explanation.

The underlying reality has not changed. Bones, muscles and tendons continue to perform multiple biological functions in both investigations. What has changed is the representation selected for the purpose of explanation.

This illustrates an important principle of General Systems Theory. Components, subsystems and systems of interest are representational roles rather than fixed ontological categories. The same entity may be represented differently in different investigations according to the explanatory relevance of its own emergent properties, behaviours or outputs.

Choosing an appropriate representation therefore becomes an essential part of systems investigation. Representing unnecessary levels of organisation obscures the principal causal mechanisms, while representing too little may omit causal processes that are essential to explanation. Effective systems investigation therefore seeks an appropriate level of representation that balances explanatory power with conceptual simplicity.

This perspective is entirely consistent with Critical Realism. The underlying causal organisation exists independently of the observer, while the representation adopted reflects the analytical purpose of the investigation. Changing the representation changes neither the underlying reality nor its causal organisation. It changes only the level of causal organisation that is explicitly represented.

6.3 The System of Interest

Physical reality exists independently of the observer. It comprises countless interacting entities, processes, and systems that together form an interconnected hierarchy. Scientific investigation does not create these systems but seeks to recognise and represent those relevant to a particular purpose.

In practice, however, it is neither possible nor desirable to investigate every aspect of reality simultaneously. Scientific enquiry therefore begins by selecting a system of interest (Checkland, 1981). The system of interest is the particular system chosen for investigation because it is considered relevant to the questions being asked.

The selection of a system of interest is therefore an epistemic act rather than an ontological one. The underlying system is assumed to exist independently of observation, but the decision to study it, the level at which it is represented, and the boundaries assigned to it depend upon the analytical purpose of the investigation.

Consider a river catchment. A hydrologist concerned with flood prediction may select the entire catchment as the system of interest. An ecologist may instead focus upon a wetland within the catchment, while a civil engineer may concentrate on the operation of a particular flood defence. Each investigation considers a different system of interest, even though all concern the same underlying physical reality.

The selection of a system of interest is closely related to the identification of a system boundary. A boundary distinguishes those entities, processes, and interactions considered to form part of the system from those regarded as lying outside it. Boundaries are therefore representational devices that assist analysis. They should not necessarily be interpreted as absolute physical discontinuities. In many cases, systems remain strongly coupled to their surroundings through continual exchanges of matter, energy, and information.

The environment of a system should not be understood simply as everything lying outside the system boundary. Rather, it comprises those entities, processes, and systems lying outside the boundary that are capable of causal interaction with the system of interest. Entities lying outside the boundary but incapable of influencing, or being influenced by, the system are generally irrelevant to its analysis and need not be considered part of its effective environment.

Systems may also be understood in terms of their relationships with other systems. A system may form part of a larger parent system, contain one or more subsystems or child systems, and interact with other sibling systems that share the same parent. These family relationships provide a convenient conceptual shorthand for describing hierarchical organisation and interactions between systems. They should not be regarded as fundamental ontological categories but as useful representational concepts that assist communication and analysis.

The choice of system of interest consequently influences every subsequent stage of analysis, including the identification of components, the representation of processes, the recognition of system properties, and the interpretation of system behaviour. Different investigators may legitimately select different systems of interest according to their objectives without implying disagreement concerning the underlying reality.

The concept of a system of interest therefore provides an important bridge between ontology and epistemology. Ontologically, systems are assumed to exist independently of observation. Epistemically, investigators select, bound, and represent those systems that are most appropriate for the questions they seek to answer. This also helps explain why different scientific disciplines often investigate different aspects or levels of the same underlying reality. Their conclusions need not be contradictory because they may simply reflect different analytical purposes and therefore different selections of the system of interest.

The definition adopted in this framework is therefore:

System of Interest: The system selected and represented for investigation according to the purpose of the enquiry.

6.4 Components

Systems are composed of components that participate in the organisation and behaviour of the system as a whole. A component is any identifiable entity, process, or subsystem that contributes to the organisation, properties, behaviours, or outputs of the system of which it forms a part.

Components should not be regarded simply as physical parts. Depending upon the analytical purpose of the investigation, a component may be represented as a physical entity, a process, a subsystem, or another identifiable system. The representation selected depends upon the analytical purpose rather than upon any fundamental difference in the underlying reality.

The relationships between components are often as important as the components themselves. Components interact through the transfer of matter, energy, or information, and it is these interactions that give rise to the organisation of the system. A collection of components without significant causal interaction does not constitute a functioning system.

Components may themselves be represented as systems (Boulding, 1956; Koestler, 1967). Such systems are commonly referred to as subsystems because they form part of a larger system. Conversely, every system may itself be regarded as a component of a larger supersystem (sometimes termed a suprasystem). Systems therefore exist within hierarchies in which each level may simultaneously be viewed as a system in its own right, as a subsystem of a larger system, and as a supersystem containing smaller systems.

The level at which components are identified depends upon the analytical purpose of the investigation. A motor vehicle may be represented as comprising an engine, transmission, suspension, and braking system. The engine may then be represented as comprising cylinders, pistons, valves, and lubrication systems. At a finer level still, those components may themselves be represented as systems of interacting mechanical parts. There is therefore no single correct level of decomposition. The appropriate level depends upon the questions being addressed.

Similarly, systems may be structured into larger systems. Individual trees may be represented as components of a forest ecosystem; individual businesses as components of an economy; and individual organs as components of an organism. Decomposition and structuring are therefore complementary representational processes that allow investigators to move between different levels of organisation while maintaining continuity with the underlying physical reality.

The concepts of component, subsystem, and supersystem should therefore be understood as relative rather than absolute. Whether an entity is regarded as a component, a subsystem, or a supersystem depends upon the system of interest selected by the observer. The same physical entity may legitimately be represented in each of these ways according to the analytical purpose of the investigation.

The definitions adopted in this framework are therefore:

Component: An identifiable entity, process, assembly or system represented as contributing to the organisation, behaviour or outputs of a larger assembly or system.

Sub-system: A system that forms part of a larger system.

Super-system: A system containing one or more sub-systems.

Parent System: A system containing one or more sub-systems.

Because systems are nested hierarchically, the distinction between a component of a system and a system is primarily one of representational perspective rather than ontology. The same causal organisation may be represented as a system in one investigation and as a subsystem in another according to the analytical purpose and the level of organisation being considered.

6.5 System Properties

Like all entities, systems possess properties that characterise their nature and behaviour. Some of these properties are inherited directly from their constituent components, while others arise from the organisation of those components within the system. Distinguishing between these different kinds of properties is important because it explains why systems frequently exhibit characteristics that cannot be attributed to their individual components considered in isolation.

For convenience, system properties may be grouped into four broad categories: component properties, organisational properties, behavioural properties, and functional properties.

Component properties originate in the constituent entities of the system. These include characteristics such as mass, size, composition, temperature, position, and other properties introduced earlier in this paper. Although these properties belong to individual components, they frequently influence the behaviour and characteristics of the system as a whole.

Organisational properties describe how a system is organised rather than the properties of its individual components. They include characteristics such as configuration, organisation, coupling, openness, and the nature of the system boundary. Together, these properties describe how the components are arranged, interact, and are represented as an organised whole.

Behavioural properties describe how a system behaves through time. These include persistence, stability, adaptability, resilience, and other characteristics associated with the continued operation of the system under changing conditions. Such properties are commonly recognised through observation of the system's behaviour and outputs rather than by direct inspection of its constituent components.

Functional properties describe what the system does. A pump transports fluid, a manufacturing system produces goods, and a biological cell maintains the processes necessary for life. Function and purpose are considered more fully later in this paper, but they are introduced here because they provide one of the principal means by which many system properties become observable.

Together, these categories distinguish the properties of individual components from those of the organised whole. While some system properties depend directly upon the characteristics of individual components, others arise from the organisation of those components within the system.

Of particular importance are those properties that arise from organisation rather than from the components considered individually. These are referred to as emergent properties and constitute one of the defining characteristics of systems. Emergent properties are generally recognised through the behaviours and outputs they enable rather than through direct observation. Their causal basis is developed in a subsequent paper and is therefore introduced only briefly here.

The classification presented in this section provides a practical framework for describing, comparing, and analysing systems while maintaining a clear distinction between the properties of individual components and those of the organised whole.

6.6 Emergent Properties

One of the defining characteristics of systems is that they exhibit emergent properties (von Bertalanffy, 1968). Within this framework, the presence of identifiable emergent properties, behaviours, or outputs distinguishes a system from a pre-systemic assembly. Emergent properties arise from the organisation of the system rather than from its individual components considered in isolation (Broad, 1925).

Some system properties are attributable directly to the properties of their individual components. The total mass of a system, for example, is determined by the masses of its constituent components. Other properties, however, depend upon the organisation of those components and the causal interactions between them. The load-bearing capability of an arch, the pumping action of the heart, the information-processing capability of a computer, and the viability of a living cell all arise from the organisation of their constituent components rather than from any individual component alone.

Emergent properties are not normally observed directly. Rather, they are recognised through the behaviours and outputs they enable. A bridge demonstrates its structural capability by supporting loads. A pump demonstrates its function by producing flow. A living organism demonstrates its organisation through growth, repair, and reproduction. It is through such recurring behaviours and outputs that the underlying emergent properties of systems are recognised.

The significance of emergence extends beyond the appearance of new properties. The emergence of new properties, behaviours, or outputs creates new causal possibilities. Once a system exhibits emergent capabilities, it may participate in causal interactions that were not possible for its constituent components acting individually. Emergence therefore increases the causal capabilities of organised systems and provides the foundation for the formation of higher levels of causal organisation.

Emergence therefore provides one of the principal characteristics distinguishing systems from simpler organised assemblies. The present paper introduces emergence only as a defining characteristic of systems. The causal processes through which emergent properties arise, create new causal possibilities, and contribute to the recursive organisation of reality are developed in a subsequent paper.

Organisational properties describe how a system is organised rather than the properties of its individual components. They include characteristics such as configuration, organisation, structure, coupling, openness, and the nature of the system boundary. Together, these properties describe how the components are arranged, interact, and are represented as an organised whole.

The definitions adopted in this framework are therefore:

Emergence: The process by which the organisation and causal interaction of components gives rise to properties, behaviours, outputs, or causal capabilities that are not exhibited by the components considered in isolation.

Emergent Property: A property arising through emergence.

Emergent Behaviour: A behaviour arising through emergence.

Emergent Output: An output arising through emergence.

6.7 Provisional Definition of a System

The preceding sections have progressively developed the concepts required to understand systems from first principles. Beginning with entities, properties, relationships, interactions, and processes, progressively richer forms of causal organisation were introduced through collections, configurations, assemblies, organisation, and structure. The concept of emergence then provided an objective criterion by which systems may be distinguished from pre-systemic assemblies.

Within this framework, a system is not defined simply by the existence of interacting components. Rather, it is defined by the emergence of identifiable properties, behaviours, or outputs arising from the organisation of those components. These emergent characteristics distinguish the organised whole from its constituent parts considered independently.

The definition also reflects the distinction maintained throughout this paper between ontology and epistemology. The causal organisation of a system is assumed to exist independently of observation. However, the recognition, delimitation, and representation of a system depend upon the analytical purpose of the observer. A system is therefore both an ontological phenomenon and an epistemic representation of that phenomenon.

The preceding sections have developed and justified the concepts contained within the working definition of a system introduced earlier. That definition emphasises that systems are organised assemblies possessing emergent properties, behaviours, or outputs.

For scientific investigation, systems are recognised and represented as wholes according to the analytical purpose of the observer. Their boundaries, components, and level of representation are therefore epistemic choices made in support of a particular investigation rather than intrinsic features of the underlying reality.

This definition is intended as a foundational definition for the present framework. Subsequent papers develop particular aspects of it in greater detail, for example, the causal mechanisms through which interactions occur and the emergence of new properties, behaviours, and causal capabilities arising from organised systems.

7. Systems, Boundaries and Environments

7.1 System Boundaries

A system boundary distinguishes the system selected for investigation from its environment. Boundaries provide an essential means of representing systems because they identify which entities, processes, interactions, and transfers are considered part of the system and which are external to it.

Boundaries should not, however, be interpreted simply as ontological discontinuities. Physical reality is generally continuous. Matter, energy, and information frequently pass across the boundaries we assign to systems, and causal interactions often extend beyond them. Open systems continually exchange matter, energy, or information with their environments (von Bertalanffy, 1968), while even apparently closed systems may interact with their surroundings to some extent.

The existence of ontological continuity does not imply that system boundaries are arbitrary. Reality contains identifiable organisations that differ in their internal organisation, patterns of interaction, and causal capabilities. These organisations provide an objective basis for recognising systems. The role of the observer is not to invent those organisations but to select boundaries that most appropriately represent them for the purpose of a particular investigation.

System boundaries are therefore epistemic representations of ontological organisations within a continuous reality. Different investigators may legitimately draw different boundaries around the same underlying organisation according to the questions they seek to answer. Nevertheless, those choices are constrained by the organisation that actually exists. Boundaries are therefore neither entirely objective nor entirely subjective. They represent an interaction between the structure of reality and the analytical purpose of the observer.

Boundaries perform more than a simple separating function. They provide the reference by which entities, components, processes, interactions, and transfers are classified as internal or external to the system of interest. Inputs and outputs are likewise defined relative to the selected boundary. The choice of boundary therefore influences the representation of the system, the identification of its components, the recognition of its properties, and the interpretation of its behaviour.

The relationship between boundaries and organisation may also be understood in terms of causal complexity. The interactions occurring within a system frequently exhibit patterns of organisation and causal integration that differ from those exhibited by its individual components or by the surrounding environment. The causal transfers associated with a living cell, for example, are considerably more numerous and more highly organised than those associated with the individual molecules from which the cell is composed. Similarly, the interactions within an organisation, an ecosystem, or a society often exhibit forms of organisation that are not evident when their individual components are considered independently. Such differences in organisation provide an important basis for recognising systems and selecting appropriate system boundaries.

The selection of a system boundary is therefore closely related to the concept of the system of interest developed in the previous section. Different boundaries may be equally appropriate for different analytical purposes, provided they correspond sufficiently with the underlying organisation being investigated.

The definition adopted in this framework is therefore:

System Boundary: The epistemic representation that distinguishes a system of interest from its environment for the purposes of investigation.

7.2 Internal and External

The terms internal and external are widely used in systems theory, but they should not be understood as absolute properties of entities or processes. Whether something is internal or external depends entirely upon the system of interest and the boundary selected to represent it.

An entity may therefore be internal to one system while simultaneously being external to another. A human heart, for example, is internal to the human body but external to an individual cell within that body. Similarly, the heart forms part of the environment of a blood cell circulating through the cardiovascular system. The same physical entity therefore participates in multiple systems and occupies different relationships within each.

The same principle applies to processes and interactions. A biochemical process occurring within a cell is internal to that cell but external to one of its constituent molecules. Likewise, the circulation of blood is internal to the cardiovascular system but external to an individual organ considered as the system of interest. Internality and externality therefore describe relationships between entities and systems rather than intrinsic characteristics of the entities themselves.

The distinction between internal and external is established by the system boundary. Entities, processes, interactions, and transfers represented as lying within the selected boundary are regarded as internal to the system. Those represented as lying outside the boundary are regarded as external. Since the selection of the boundary depends upon the analytical purpose of the investigation, the distinction between internal and external is likewise relative to that purpose.

This relativity should not be interpreted as implying that the distinction is arbitrary. The system boundary is selected to represent an underlying organisation within physical reality, and the classification of entities and processes as internal or external is constrained by that organisation. Different investigators may legitimately classify the same entity differently because they have selected different systems of interest, not because reality itself has changed.

Recognising the relativity of internal and external relationships is particularly important when comparing systems descriptions from different disciplines. Apparent disagreements may arise simply because investigators have selected different systems of interest and therefore different system boundaries. Once these differences are recognised, apparently conflicting descriptions frequently prove to be complementary representations of the same underlying reality.

The definitions adopted in this framework are therefore:

Internal: Situated within the boundary of the selected system of interest.

External: Situated outside the boundary of the selected system of interest.

7.3 Environment

Every system exists within an environment comprising the external causal networks with which it is capable of interacting (von Bertalanffy, 1968). The environment provides the wider context within which the system operates and with which it may exchange matter, energy, or information. The concept of environment is therefore fundamental to the analysis of open systems.

The environment should not be understood simply as everything lying outside the system boundary. Rather it comprises all external causal networks capable of influencing the system or being influenced by it. These causal networks may be assemblies, pre-systemic assemblies or complete systems. They may provide resources, impose constraints, exchange matter, energy or information, or otherwise influence the behaviour of the system.

Assemblies frequently form parts of larger systems. Consequently, many environmental causal networks may themselves be represented as systems hierarchies. A system of interest may therefore interact simultaneously with multiple external systems, each of which may form part of its own recursively nested hierarchy.

For example, a surgeon interacts directly with an operating theatre, but also with family, transport systems, professional organisations, government regulation and wider society. Each of these constitutes an external causal network capable of interacting with the surgeon. The operating theatre forms part of a surgical department, which forms part of a hospital, which forms part of a national healthcare system. Likewise, the surgeon's family, professional organisations and transport systems each belong to different systems hierarchies.

For the purposes of investigation, these interacting causal networks are often represented by one or more environmental hierarchies. Closely coupled systems typically obtain most of the resources necessary for their functioning from one dominant parent system and are therefore conveniently represented by a single environmental hierarchy. More loosely coupled systems obtain important resources from several external systems simultaneously and therefore participate in several overlapping environmental hierarchies.

The extent and nature of the environmental representation depends upon the analytical purpose of the investigation. A detailed study of a biological cell may focus upon the surrounding tissue as its principal environment, whereas a physiological investigation may instead consider the environment of the entire organism. Similarly, the environment of an organisation may comprise customers, suppliers, regulators, competitors, and other organisations with which it interacts, while a national economy may itself form part of the environment of that organisation in a broader investigation.

The concept of environment therefore illustrates once again the relationship between ontology and epistemology developed throughout this paper. Ontologically, countless entities, processes, and systems coexist within a continuous reality. Epistemically, investigators identify those external entities and processes that are causally relevant to the system of interest. The environment is therefore not an arbitrary remainder left after defining the system, but a representation of those parts of reality that participate in significant causal interactions with it.

The definition adopted in this framework is therefore:

Environment: The set of external causal networks, including assemblies and systems, capable of causal interaction with a system of interest.

7.4 Component Environments

Just as systems are organised hierarchically, so too are the environments experienced by their components. A component's environment comprises the external causal networks capable of interacting with it. These include its parent system together with any other external causal networks from which it receives, or to which it transfers, matter, energy or information.

For many components, particularly those that are closely coupled to their parent system, the parent system provides most of the resources, constraints and interactions necessary for continued functioning. A heart, for example, derives most of its matter, energy and information from the body and therefore experiences the body as its dominant immediate environment.

Other components interact much more extensively with causal networks lying beyond their parent system. A leaf obtains sunlight, atmospheric gases and rainfall directly from the wider environment, while a sales department interacts directly with customers, suppliers and competitors. Such components therefore participate simultaneously in several environmental causal networks.

Many of these environmental causal networks are themselves organised as larger systems. Consequently, a component may participate simultaneously in several environmental hierarchies, each representing a different aspect of its interactions with the wider world. Which of these hierarchies is represented depends upon the analytical purpose of the investigation.

The degree to which a component depends upon its parent system rather than upon other external causal networks is described by its degree of coupling. This relationship is examined in the following section.

The definitions adopted in this framework are therefore:

Component Environment: The external causal networks capable of interacting with a component.

Parent Environment: The environment of the component's parent system within which the parent system itself interacts.

7.5 Coupling

The environments experienced by a subsystem depend not only upon the external causal networks with which it interacts but also upon the degree to which it depends upon each of those networks for its continued functioning. This degree of dependence is referred to as coupling.

A subsystem may interact with many external causal networks simultaneously. Some of these interactions provide resources essential for continued functioning, while others have only occasional or minor influence. Coupling describes the relative importance of these interactions rather than simply their existence.

A closely coupled subsystem derives most of the matter, energy and information necessary for its continued functioning from a parent system. That parent system therefore provides the subsystem's dominant environment and mediates many of its interactions with the wider world.

A more loosely coupled subsystem derives a greater proportion of the matter, energy and information necessary for its continued functioning directly from other external causal networks beyond a parent system. Such subsystems therefore participate more extensively in multiple overlapping environments and are influenced more directly by changes occurring outside the parent system.

The distinction is illustrated by comparing different biological and social systems. A heart depends primarily upon interactions occurring within the body and is therefore closely coupled to that parent system. By contrast, the leaves of a tree obtain sunlight, atmospheric gases and rainfall directly from the wider environment, making them more loosely coupled to the tree. Similarly, a production department may depend primarily upon internal organisational processes, whereas a sales department interacts extensively with customers, suppliers and competitors beyond the organisation itself.

Coupling therefore determines the relative importance of the environmental causal networks within which a subsystem operates. Closely coupled subsystems are often represented satisfactorily by a single dominant environmental hierarchy because most significant interactions are mediated through a parent system. More loosely coupled subsystems obtain essential inputs from several environmental causal networks simultaneously and are therefore more appropriately represented by several overlapping environmental hierarchies. The hierarchy or hierarchies represented depend upon the analytical purpose of the investigation.

Coupling should therefore be understood not simply as the strength of interaction between systems, but as the degree to which the continued functioning of a subsystem depends upon causal inputs originating within a parent system rather than from other external causal networks. It therefore provides an important link between system organisation, openness, environmental interaction, hierarchy and the representation of systems for investigation.

The definitions adopted in this framework are therefore:

Coupling: The degree to which the continued functioning of a subsystem depends upon causal inputs originating within a parent system rather than from other external causal networks.

Close or Tight Coupling: A high degree of dependence upon causal inputs generated within a parent system.

Loose Coupling: A relatively low degree of dependence upon causal inputs generated within a parent system, with greater dependence upon other external causal networks.

7.6 Perspective and Boundary Selection

The concepts developed throughout this paper demonstrate that the representation of systems depends not only upon the organisation present in physical reality but also upon the analytical purpose of the investigation. Different investigators may therefore legitimately select different systems of interest, boundaries, levels of organisation, components, and environments according to the questions they seek to answer.

These differences should not be interpreted as implying that different investigators inhabit different realities. The underlying physical reality remains the same. What differs is the aspect of that reality selected for investigation and the representational framework used to describe it.

A hydrologist may represent an entire river catchment as the system of interest, while an ecologist may focus upon a wetland within the catchment, and a civil engineer upon a flood defence structure. Each investigator selects different boundaries, components, environments, and variables because each seeks to understand different aspects of the same underlying reality. The resulting representations are therefore complementary rather than necessarily contradictory.

Perspective consequently influences every stage of systems representation. It affects the selection of properties, the identification of components, the recognition of organisation and structure, the placement of boundaries, the definition of environments, and the level of organisation at which systems are represented. None of these choices is arbitrary, because each is constrained by the underlying causal organisation of reality. Equally, none is uniquely determined, because different purposes legitimately require different representations.

This relationship between reality and representation provides an important foundation for poly-perspectivism. Different perspectives frequently describe overlapping aspects of the same reality rather than incompatible realities. Apparent disagreements may therefore arise because investigators have selected different systems of interest, different levels of organisation, or different representational frameworks rather than because they disagree about the underlying phenomena.

Recognising this distinction is an important prerequisite for productive co-ordination. Before attempting to compare theories or reconcile terminology, it is first necessary to understand the systems of interest, boundaries, levels of organisation, and purposes that underlie each representation. Once these are recognised, many apparent disagreements prove to reflect differences of perspective rather than genuine contradictions.

The principles developed in this paper therefore support a systematic approach to comparing and integrating alternative systems representations. They provide a conceptual foundation for the methods of productive co-ordination developed in the author's earlier paper of that name (Challoner, 2026).

8. Inputs, Processes, Outputs and Representation

8.1 Inputs

Systems interact with their environments through the transfer of matter, energy, and information. Transfers entering the system are commonly represented as inputs, while those leaving the system are represented as outputs (von Bertalanffy, 1968).

An input should not be understood simply as something received by a system. More generally, it is anything that becomes available to participate in the causal organisation of a system. Whether a particular entity, process, or transfer is represented as an input therefore depends upon the system selected for investigation and the analytical purpose of the representation.

Inputs may comprise matter, energy, information, or combinations of these. A plant receives water, carbon dioxide, mineral nutrients, and solar energy as inputs. A manufacturing system receives raw materials, energy, and design information. A living organism receives food, oxygen, sensory information, and many other inputs from its environment. Similarly, a social organisation may receive people, materials, financial resources, legislation, and information from other organisations or from wider society.

The representation of inputs is epistemic rather than ontological. The entity, process, or transfer itself exists independently of observation, but its classification as an input depends upon the system selected for investigation and the boundary used to represent that system. An entity represented as an output from one system may simultaneously be represented as an input to another system. It may also become incorporated into a larger system, where it functions as one of that system's internal components or participates in one of its internal processes.

Inputs should not be regarded as isolated acquisitions. They frequently participate in organised networks of causal interaction through which they are transformed, stored, combined, or redistributed within the system. Their significance therefore depends not only upon their individual characteristics but also upon the organisation of the internal processes with which they interact.

For the purposes of general systems theory, inputs may frequently be represented more simply as transfers that cross the boundary of the receiving system from its environment and become incorporated into its internal causal organisation. This representation is sufficient for most systems analyses and provides the basis for the familiar Input–Process–Output (IPO) representation introduced later in this section. More general situations involving the incorporation of components, assemblies, or complete systems are accommodated by the broader definition adopted here.

The definition adopted in this framework is therefore:

Input: Matter, energy, information, components, assemblies, systems, or processes that become available to participate in the causal organisation of a system.

8.2 Outputs

Systems influence their environments through the transfer of matter, energy, and information across their boundaries (von Bertalanffy, 1968). Transfers leaving the system are commonly represented as outputs, while those entering the system are represented as inputs.

An output should not be understood simply as something produced by a system. More generally, it is anything that becomes independently available to participate in causal interactions beyond the internal organisation from which it arises. Whether a particular entity, process, or transfer is represented as an output therefore depends upon the system selected for investigation and the analytical purpose of the representation.

Outputs may comprise matter, energy, information, or combinations of these. A plant releases oxygen, water vapour, seeds, and thermal energy to its environment. A manufacturing system produces finished goods, waste materials, heat, and information concerning production. Living organisms produce movement, sound, heat, waste products, and many other outputs. Similarly, social organisations may produce goods, services, financial transactions, decisions, legislation, knowledge, or other forms of information that influence their wider environments.

As with inputs, the representation of transfers as outputs is epistemic rather than ontological. The transfer itself exists independently of observation, but its classification as an output depends upon the system selected for investigation and the boundary used to represent that system. A transfer represented as an output from one system may simultaneously be represented as an input to another system or as an internal transfer within a larger system.

Outputs should not be regarded as isolated products. They frequently participate in wider networks of causal interaction, becoming inputs to other systems or contributing to processes occurring within the environment. Through these continuing causal interactions, outputs may influence the behaviour of other systems, modify environmental conditions, or participate in the emergence of larger-scale organisations and processes.

The entities represented as outputs need not be limited to simple transfers of matter, energy, or information. They may themselves constitute organised assemblies or complete systems. An atom, for example, is a complete system in its own right but may simultaneously serve as an input to the assembly of a molecule. Similarly, a manufactured engine is itself a system but may be represented as the output of one manufacturing system and the input to the assembly of a motor vehicle. Conversely, outputs may arise through the disassembly of an existing system, producing multiple components, assemblies, or simpler transfers that subsequently participate in other systems. Whether an entity is represented as a system, a component, an input, or an output therefore depends upon its role within the causal organisation under investigation and the analytical purpose of the representation.

For the purposes of general systems theory, outputs may frequently be represented more simply as transfers that leave the boundary of the system of origin and become available for causal interaction externally. This representation is sufficient for most systems analyses and provides the basis for the familiar Input–Process–Output (IPO) representation introduced later in this section. More general situations involving the assembly, disassembly, division, or fusion of systems may require changes in the representation of system boundaries. These aspects of system formation, transformation, and life cycles are developed further in a subsequent paper on Assembly Theory and the Formation of Complex Systems.

The recognition of outputs also provides an important basis for identifying emergent properties. Many emergent properties become evident through the behaviours and outputs exhibited by systems rather than through direct observation of their internal organisation. Outputs therefore play an important role not only in the causal interaction between systems but also in the recognition and representation of systems themselves.

The definition adopted in this framework is therefore:

Output: Matter, energy, information, components, assemblies, systems, or processes that become independently available to participate in causal interactions beyond the internal organisation from which they arise.

8.3 Internal Processes

An internal process is not a different kind of process but a process considered with respect to a selected system boundary. Processes represented as lying within the boundary of the system of interest are described as internal processes, while those represented as lying outside the boundary can form part of the system's environment.

Internal processes receive inputs from other processes within the system, from the system's external environment, or from a combination of both. Through the organised interactions of the system's components, these inputs are transformed, stored, redistributed, or otherwise participate in further causal interactions. The resulting outputs may remain within the system as inputs to other internal processes or become independently available to participate in causal interactions beyond the system.

The organisation of internal processes distinguishes systems from simple collections of interacting entities. A system is recognised not merely because processes occur within it, but because those processes are causally organised so that their interactions collectively contribute to the behaviour of the system as a whole. Individual internal processes may therefore be understood as participating in a wider causal organisation whose properties, behaviours, and outputs may differ from those of the processes considered independently.

Internal processes frequently interact in networks rather than simple linear sequences. The outputs of one process may provide inputs to several others, while individual processes may receive inputs from multiple sources. Such networks often include feedback, branching, convergence, storage, and other forms of causal organisation that contribute to the behaviour of the system. These more complex forms of organisation are considered in greater detail in later papers but should be recognised here as common features of many systems.

The distinction between internal and external processes is epistemic rather than ontological. A process represented as internal to one system may be represented as an independent system or as part of the environment of another according to the system of interest and the analytical purpose. The underlying causal processes remain unchanged; only their representation differs.

The definition adopted in this framework is therefore:

Internal Process: A process represented as occurring within the boundary of a selected system and participating in its causal organisation.

8.4 The Input–Process–Output Representation

The concepts of inputs, internal processes, and outputs provide the basis for one of the most widely used representations in systems theory (Hall & Fagen, 1956). A system is represented as receiving one or more inputs from its environment, transforming those inputs through internal processes, and producing one or more outputs that subsequently participate in wider causal interactions.

In its simplest form, this representation may be expressed as:

Inputs → Processes → Outputs

The Input–Process–Output (IPO) representation should not be understood as the system itself. Rather, it is a simplified representation that foregrounds the principal causal transformations occurring within the system while suppressing much of the underlying organisational detail. Networks of interacting processes, feedback loops, storage, branching pathways, and many other aspects of causal organisation are compressed into a representation that is readily understood and communicated.

The IPO representation therefore constitutes a form of causal compression. Rather than representing every individual entity, interaction, and process, it represents the principal causal relationships relevant to the analytical purpose. This simplification allows systems to be compared, analysed, communicated, and reasoned about without requiring the full complexity of their internal organisation to be represented.

The same representation may be interpreted in terms of the Transfer–Process–Transfer (TPT) representation introduced in a subsequent paper. Inputs and outputs correspond to transfers of matter, energy, or information across the system boundary, while the internal processes represent the causal transformations through which those transfers are organised. The IPO and TPT representations therefore describe the same underlying causal organisation using different but complementary representational idioms.

Like all representations considered throughout this framework, the IPO representation is epistemic rather than ontological. The underlying causal organisation exists independently of its representation, while the IPO representation provides a purposeful simplification selected according to the needs of the investigation. Different analytical purposes may therefore justify different levels of causal compression and different representations of the same underlying system.

The definition adopted in this framework is therefore:

Input–Process–Output (IPO) Representation: A simplified representation of a system that describes its causal organisation in terms of inputs, internal processes, and outputs, providing a purposeful causal compression of the underlying organisation.

8.5 Strengths and Limitations of the IPO Representation

The Input–Process–Output (IPO) representation provides a simple and widely applicable means of describing systems. By representing systems in terms of inputs, internal processes, and outputs, it captures the principal causal transformations occurring within a system while avoiding unnecessary detail. This simplicity makes the representation particularly valuable for communication, comparison, analysis, education, and the preliminary investigation of systems.

The principal strength of the IPO representation lies in its ability to compress complex causal organisations into a form that is readily understood. Many engineering, biological, organisational, and social systems may be usefully represented in this way because attention is focused upon the principal transfers and transformations relevant to a particular analytical purpose.

The usefulness of the IPO representation, however, arises from the very simplifications that also limit it. As a form of causal compression, it suppresses much of the underlying organisation responsible for the behaviour of the system. Internal networks of interacting processes, multiple pathways of causal influence, storage, branching, convergence, and other aspects of causal organisation are represented only indirectly or omitted entirely.

The IPO representation also provides only a limited account of feedback. Feedback loops, recursive interactions, and dynamic patterns of mutual influence may often be represented only by extending or modifying the basic IPO representation. Similarly, systems comprising many interacting subsystems may require network representations that describe the causal organisation more explicitly.

The representation is likewise influenced by the selection of the system of interest. Different investigators may legitimately choose different system boundaries, levels of organisation, or analytical purposes, resulting in different but equally valid IPO representations of the same underlying reality. The representation therefore reflects both the causal organisation of the system and the purpose for which that organisation is being represented.

The IPO representation should also not be regarded as a complete description of emergence. While emergent properties are frequently recognised through the behaviours and outputs represented within an IPO framework, the representation itself does not explain how those properties arise from the organisation of interacting components and processes. The causal basis of emergence is considered separately in a subsequent paper.

These limitations should not be regarded as deficiencies of the IPO representation but as natural consequences of purposeful representational compression. Every representation preserves those aspects of reality relevant to its intended purpose while suppressing others. The IPO representation therefore remains an extremely valuable systems idiom, provided its strengths, limitations, and intended purpose are properly understood.

8.6 Historical Note – Throughput

The term throughput has been widely used in systems theory to describe what occurs between the inputs and outputs of a system (von Bertalanffy, 1968). In different contexts it has referred to the internal transformations occurring within a system, the movement of matter, energy, or information through the system, or the overall processing of inputs into outputs.

Within the present framework these ideas are represented more explicitly. Internal transformations are described as processes, while the movement of matter, energy, and information is represented in terms of transfers. Where appropriate, the rates at which such transfers or processes occur may be represented as properties of the system rather than as distinct systems concepts.

The term *throughput* may therefore remain useful as a convenient descriptive expression when referring informally to the overall passage and transformation of matter, energy, or information through a system. However, it is not adopted as a fundamental concept within the systems idiom developed in this framework because its various meanings are more precisely represented by the concepts of transfers, processes, and their associated properties.

Readers familiar with the systems literature will therefore recognise *throughput* as a historical term that broadly corresponds to the internal causal activity represented by the Input–Process–Output representation, but the present framework prefers the more specific terminology introduced in the preceding sections.

9. System States and Behaviour

9.1. Introduction

Systems are not simply organised collections of components. They also exhibit states and behaviours that change through time. The concepts introduced in this section provide progressively more compressed representations of the changing condition of a system. Individual variables describe selected properties of a system. Their values collectively represent the system's state at a particular time. Successive states provide the basis for describing system behaviour, while recurring patterns of behaviour may be represented more generally as system dynamics. Each stage preserves those relationships relevant to its intended analytical purpose while reducing representational complexity.

9.2 System State

The properties of a system may be represented by selected variables whose values describe particular aspects of the system at a given time. Collectively, these values provide a representation of the condition of the system. This representation is known as the system state (von Bertalanffy, 1968).

A system state is not the system itself but an epistemic representation of selected aspects of the system at a particular time. It is obtained by combining the values of those variables considered relevant to the analytical purpose. Different investigations may therefore represent the same underlying system using different sets of variables and consequently different state representations.

The concept of system state provides a further stage of representational compression. Individual values describe selected properties of the system. The collection of these values is subsequently compressed into a single representation of the system's condition at a specified time. This enables complex systems to be compared, analysed, and reasoned about without requiring every property of the system to be represented explicitly.

The choice of variables depends upon the analytical purpose of the investigation. The state of a river may be represented by variables such as water level, flow rate, sediment load, and water quality. The state of a manufacturing system may be represented by production rate, inventory, energy consumption, and machine availability. Similarly, the state of a social system may be represented by selected demographic, economic, institutional, or cultural variables. No representation captures every property of the underlying system; each represents only those considered relevant to the purpose of the analysis.

The system state should therefore not be confused with the complete condition of the system. It is a purposeful representation constructed from selected variables and values. The validity of the representation depends upon the extent to which the selected variables preserve those aspects of the underlying system relevant to the intended analytical purpose.

The definition adopted in this framework is therefore:

System State: A representation of the condition of a system at a specified time obtained from the values of selected system variables.

The concept of system state provides the basis for describing how systems change through time. Successive states allow patterns of change to be recognised and represented, leading naturally to the concepts of state change and system behaviour developed in the following sections.

9.3 System State Change

A system state represents the condition of a system at a specified time. As the properties of the system change, the values of one or more system variables may also change. The resulting change in the representation of the system is described as a state change (von Bertalanffy, 1968).

A state change is not itself a physical process but an epistemic representation of changes occurring within the underlying system. The causal interactions, processes, and transfers that occur within the system give rise to changes in its properties. These changes are subsequently represented as changes in the values of the selected system variables and, consequently, as changes in the system state.

State changes may involve one or many variables. Some changes may be small and continuous, while others may be abrupt or discontinuous. The significance of a state change depends not only upon the magnitude of the changes in individual variables but also upon the analytical purpose and the aspects of the system represented by those variables.

The transition from one state to another provides a convenient means of representing the evolution of a system through time. Rather than describing every individual causal interaction occurring within the system, successive states provide a compressed representation of its changing condition. The sequence of these state representations enables systems to be compared, monitored, and analysed while suppressing much of the underlying causal complexity.

The representation of state change therefore provides another stage of epistemic compression. Complex networks of interacting processes are represented by changes in selected variables, while those changes are subsequently represented as transitions between successive system states. This compression allows attention to be focused upon the changing condition of the system rather than upon every individual causal interaction responsible for those changes.

The definition adopted in this framework is therefore:

System State Change: A change in the representation of a system resulting from changes in the values of one or more selected system variables.

State changes provide the basis for recognising patterns of system behaviour. While a single state change represents one transition between successive states, recurring sequences of state changes enable the behaviour of the system to be identified and analysed.

9.4 System Behaviour

Successive state changes provide a representation of how a system changes through time. Where these changes exhibit identifiable patterns, they may be represented collectively as the behaviour of the system (Ashby, 1956).

Behaviour is therefore not a property of a single system state but a representation of patterns of state change occurring over a period of time. It describes how the condition of a system evolves rather than the condition of the system at any particular instant.

The concept of behaviour provides a further stage of epistemic compression. Individual state changes are represented collectively as recognisable patterns that characterise the activity of the system. Rather than describing every successive state individually, behaviour summarises the overall manner in which the system changes through time. This allows complex temporal patterns to be recognised, communicated, compared, and analysed without representing every intermediate state explicitly.

The behaviour of a system may be represented in many ways according to the analytical purpose. A river may exhibit steady flow, periodic flooding, or seasonal variation. A biological organism may exhibit growth, reproduction, migration, or adaptation. A manufacturing system may exhibit continuous production, intermittent operation, or declining performance. Similarly, a social system may exhibit economic growth, demographic change, institutional stability, or social unrest. Each of these descriptions represents characteristic patterns of state change rather than individual states themselves.

Different observers may legitimately represent the behaviour of the same system in different ways because they select different variables, periods of observation, or analytical purposes. Behaviour is therefore not simply observed but interpreted through purposeful epistemic representation.

Behaviour should not be confused with the causal mechanisms responsible for it. Similar patterns of behaviour may arise from different causal organisations, while similar causal organisations may exhibit different behaviours under different conditions. Behaviour therefore provides an important means of recognising and comparing systems without necessarily explaining the mechanisms through which that behaviour is produced.

The definition adopted in this framework is therefore:

Behaviour: A representation of the characteristic patterns of state change exhibited by a system through time.

The concept of behaviour provides the foundation for describing the longer-term evolution of systems. Broader patterns such as cycles, trajectories, stability, and persistence build upon the representation of behaviour and are considered in the following sections.

9.5 System Dynamics

The behaviour of a system may itself exhibit recognisable longer-term patterns. These broader patterns are represented by the concept of system dynamics, which concerns the evolution of system behaviour through time (Forrester, 1961).

System dynamics therefore provides a more general representation than individual behaviours or state changes. While behaviour describes characteristic patterns of state change, system dynamics represents the trajectories, recurring patterns, and longer-term evolution of those behaviours. It provides a means of describing how systems develop, adapt, stabilise, oscillate, or otherwise evolve over extended periods.

The concept of system dynamics provides a further stage of epistemic compression. Rather than representing individual state changes or individual behaviours, it represents broader patterns that characterise the temporal evolution of the system. These patterns may include steady development, cyclical behaviour, oscillation, growth, decline, adaptation, or more complex trajectories through successive system states.

The representation adopted depends upon the analytical purpose. An ecologist may study the long-term dynamics of an ecosystem. An engineer may examine the fatigue response of a structure over time due to changing loads. An economist may investigate long-term economic trends, while a sociologist may examine demographic or institutional change. Although the variables selected and the mechanisms involved differ, each represents broader patterns of system behaviour through time.

System dynamics should not be confused with the causal mechanisms responsible for those patterns. Similar dynamic behaviours may arise from different causal organisations, while similar organisations may exhibit different dynamics under changing environmental conditions. The concept therefore provides a useful means of comparing systems and identifying characteristic patterns without requiring every underlying causal interaction to be represented explicitly.

The definition adopted in this framework is therefore:

System Dynamics: A representation of the trajectories and longer-term patterns of system behaviour through time.

The concepts introduced in this section provide a foundation for later discussions of stability, persistence, adaptation, and viability, where the long-term behaviour of systems becomes central to their analysis.

9.6 Stability and Persistence

Systems rarely remain in exactly the same state. Most continue to change through time while maintaining their overall organisation and identity. The concepts of stability and persistence provide a means of describing these longer-term characteristics of system behaviour.

Stability concerns the extent to which a system maintains or returns to characteristic patterns of behaviour following internal or external disturbance (Ashby, 1956). Some systems exhibit relatively little variation, while others undergo substantial fluctuations before returning to familiar patterns of behaviour. Stability therefore describes the behaviour of a system rather than the absence of change.

Persistence concerns the continued existence of the system as an identifiable causal organisation. A system may undergo considerable changes in its state, behaviour, components, or interactions while remaining recognisably the same system. Conversely, a system may lose its organisation and cease to persist, even if many of its individual components remain.

Both concepts therefore depend upon the analytical purpose and the level of organisation represented. What constitutes persistence for one investigation may represent transformation or replacement for another. Similarly, the degree of variation considered compatible with stability depends upon the variables selected and the purpose of the analysis.

Stability and persistence should not be regarded as absolute properties but as representations of the longer-term behaviour of systems. They provide useful means of comparing systems and describing their temporal characteristics while remaining independent of the detailed causal mechanisms responsible for those characteristics.

The concepts introduced here provide a foundation for the later development of system viability. Subsequent papers examine how systems maintain their organisation, adapt to changing circumstances, and continue to function despite internal variation and external disturbance.

The definitions adopted in this framework are therefore:

Stability: A representation of the tendency of a system to maintain or return to characteristic patterns of behaviour following disturbance.

Persistence: A representation of the continued existence of an identifiable causal organisation through time.

10. Open and Closed Systems

10.1 Open Systems

Systems exist within wider environments from which they may receive causal inputs and to which they may contribute causal outputs. A system that exchanges matter, energy, information, components, assemblies, systems, or processes with its environment is described as an open system.

Openness is therefore an ontological characteristic of a system rather than a property of its representation. It describes the existence of causal interactions between the system and its environment, irrespective of whether those interactions are represented explicitly within a particular analysis.

Open systems continually participate in wider causal networks. Matter, energy, information, components, and other causal organisations may enter the system as inputs, while outputs become independently available to participate in further causal interactions beyond the system from which they arise. These exchanges enable systems to assemble, develop, adapt, maintain their organisation, and participate in more complex systems.

The degree and significance of openness depend upon the analytical purpose. Some investigations may focus upon material exchanges, others upon energy transfers, information flows, or the movement of components and assemblies. The same system may therefore be represented as open with respect to some forms of interaction while particular exchanges are omitted for the purposes of analysis. Such omissions do not alter the openness of the underlying system but merely simplify its representation.

The recognition of open systems represented a major development in General Systems Theory. In contrast to earlier approaches that often treated systems as isolated entities, Bertalanffy emphasised that many biological and social systems depend upon continuing exchanges with their environments. The present framework extends this principle by recognising openness as a general property of causal organisations rather than one restricted to particular classes of system.

The definition adopted in this framework is therefore:

Open System: A system that participates in causal interactions with its environment through the exchange of matter, energy, information, components, assemblies, systems, or processes.

10.2 Closed and Isolated Systems

A closed system is one in which exchanges with the environment are absent or severely restricted. In thermodynamics, the term is commonly used to describe systems that exchange energy but not matter with their surroundings, while an isolated system exchanges neither matter nor energy. These definitions are appropriate within their respective scientific contexts and provide useful idealisations for the analysis of many physical systems.

From the perspective of General Systems Theory, however, complete closure is uncommon. Most biological, ecological, engineered, organisational, and social systems participate continuously in causal interactions with their environments. Consequently, the concept of a closed system is generally understood as representing an approximation in which particular forms of interaction are absent, negligible, or deliberately excluded for the purposes of analysis.

The degree of closure depends upon the forms of causal interaction being considered. A system may be closed to the transfer of matter while remaining open to the transfer of energy and the information embodied within it. A physically closed system, however, exchanges neither matter, energy, nor information with its environment. Closure should therefore be understood as relative to particular forms of causal interaction rather than as an absolute property applying equally to every aspect of a system.

The concept of closed systems nevertheless remains valuable. By reducing or eliminating particular environmental interactions, simpler causal relationships may be identified and analysed. Closed-system representations therefore provide an important foundation for many areas of science and engineering, while recognising that most real systems continue to participate within wider causal networks.

The definition adopted in this framework is therefore:

Closed System: A system in which one or more forms of causal interaction with the environment are absent or sufficiently restricted for the purposes of analysis.

Isolated System: A system that exchanges neither matter, energy nor information with its environment.

10.3 Analytical and Experimental Closure

The openness or closure of a system should not be confused with the openness or closure of its representation. In many scientific, engineering, and systems investigations, the analyst deliberately excludes selected interactions in order to simplify the investigation and focus attention upon those causal relationships considered most relevant. This deliberate simplification is referred to as analytical closure.

Analytical closure is an epistemic rather than an ontological concept. It does not alter the underlying system or eliminate the causal interactions that exist within reality. Rather, it defines the scope of the representation by specifying which interactions are included and which are deliberately omitted for the purposes of analysis.

Analytical closure is widely employed throughout science and engineering. A hydraulic model may neglect groundwater interactions when analysing river flows. A structural engineer may ignore minor thermal expansion when assessing the strength of a bridge. An economist may temporarily treat international markets as fixed while investigating a national economy. In each case, the omitted interactions continue to exist, but they are excluded from the representation because they are judged to be of secondary importance for the question under investigation.

Scientific experiments frequently seek to establish approximate physical closure by controlling or suppressing interactions with the wider environment. Such experimental closure assists in identifying particular causal mechanisms but should not be confused with analytical closure. Experimental closure concerns the conditions under which observations are made, whereas analytical closure concerns the representational decision regarding which entities and interactions are included within an explanation. In scientific research, experimental closure often facilitates analytical closure by reducing the number of interactions that need to be considered, although the two remain conceptually distinct.

Analytical closure is rarely absolute. As understanding develops, interactions that were previously omitted may be incorporated into the representation if they prove to be causally significant. Representations therefore evolve alongside improvements in scientific understanding and changes in analytical purpose.

Analytical closure therefore represents another form of purposeful epistemic compression. By excluding interactions considered irrelevant or of secondary importance, a simpler representation of the system can be developed without claiming that the omitted interactions do not exist. The validity of such representations depends upon whether the chosen analytical closure preserves the causal relationships necessary for the intended purpose.

The concepts of system boundary and analytical closure are closely related. A system boundary provides an epistemic representation of an underlying ontological continuity by identifying the system of interest. Analytical closure provides an epistemic representation of the underlying openness of that system by specifying which causal interactions are included within the representation and which are deliberately omitted. Together they define the scope of the representation without implying that reality itself has been divided or causally closed.

The definitions adopted in this framework are therefore:

Experimental Closure: The deliberate restriction or control of causal interactions between a system and its environment in order to investigate selected causal mechanisms under defined conditions.

Analytical Closure: The deliberate restriction of the causal interactions represented within an explanation in order to simplify analysis while preserving those relationships considered relevant to the analytical purpose.

10.4 Relative Openness

The distinction between open and closed systems should not be interpreted as an absolute classification. Most real systems participate in some forms of environmental interaction while restricting or excluding others. Openness is therefore best understood as a relative rather than an absolute characteristic of systems.

A system may exhibit different degrees of openness with respect to different forms of causal interaction. It may exchange matter while restricting information, exchange energy while retaining its material organisation, or exchange information extensively while remaining physically enclosed. The openness of a system therefore depends not only upon the existence of a boundary but also upon the nature of the causal interactions occurring across that boundary.

The degree of openness represented within an investigation also depends upon the analytical purpose. A representation may focus upon particular forms of exchange while deliberately omitting others through analytical closure. Consequently, the openness represented within a model or systems description should not be assumed to correspond exactly to the full openness of the underlying system.

The relative openness of systems frequently changes through time. As systems assemble, develop, adapt, or reorganise, the character and significance of their interactions with the wider environment may also change. Changes in openness may therefore influence the behaviour, persistence, and adaptability of systems and may contribute to the emergence of new forms of organisation.

The recognition of relative openness reinforces one of the central themes of this framework. The openness of a system is an ontological characteristic arising from its causal interactions with the environment, whereas the openness represented within a systems analysis depends upon the analytical purpose and the representational choices made by the investigator.

The definition adopted in this framework is therefore:

Relative Openness: The extent and character of the causal interactions between a system and its environment with respect to particular forms of exchange and a specified analytical purpose.

The relative openness of a parent system tends to influence the degree of coupling among its components. As a parent system becomes increasingly open to interactions with its wider environment, a greater proportion of the causal inputs required by its components may originate externally rather than from other components within the parent system. This generally leads to looser coupling among the components of the parent system. Conversely, where most causal inputs are generated internally, the components tend to become more tightly coupled. Relative openness and coupling therefore describe complementary aspects of the organisation of systems.

In many complex systems, interactions with the environment are not distributed uniformly across all components. Instead, they are frequently concentrated within specialised structures or processes that mediate exchanges with the wider environment. Such arrangements influence both the relative openness of the system and the coupling of its components, but their organisation and significance are considered in later papers.

10.5 Section Summary

The concepts developed in this chapter distinguish between the openness of systems in reality and the openness of their representations. This distinction allows traditional concepts such as open and closed systems to be interpreted within a broader Critical

Realist framework in which ontology and representation remain clearly distinguished. Subsequent chapters build upon these foundations when considering feedback, control, adaptation, and viability.

11. Function, Outcomes and Purpose

11.1 Outputs and Outcomes

The concept of an output was introduced earlier as anything that becomes independently available to participate in causal interactions beyond the internal organisation from which it arises. Outputs therefore represent what a system makes available to the wider causal network.

Outputs should not be confused with outcomes. An outcome is not something that leaves a system but a resulting state or condition arising from the participation of one or more outputs in subsequent causal interactions. Whereas outputs are immediately available for further interaction, outcomes describe the consequences of those interactions.

The distinction may be illustrated by a simple example. A manufacturing system produces a motor vehicle as an output. The transport of people and goods made possible by that vehicle represents an outcome. Similarly, the output of a power station is electrical energy, while the activities enabled by the use of that energy constitute outcomes. In biological systems, seeds produced by a plant are outputs, whereas the establishment of new plants is an outcome arising from their subsequent interactions with suitable environmental conditions.

Outcomes frequently extend beyond the immediate environment of the originating system. An output may participate in many successive causal interactions, giving rise to chains of outcomes occurring across multiple systems and over extended periods of time. The distinction between outputs and outcomes therefore assists in tracing causal relationships without attributing every subsequent consequence directly to the originating system.

The concepts of outputs and outcomes also represent different levels of epistemic compression. Outputs identify what becomes available for subsequent causal interaction. Outcomes represent selected consequences arising from those interactions. Which outcomes are represented depends upon the analytical purpose and the temporal and organisational scope of the investigation.

The definitions adopted in this framework are therefore:

Output: Matter, energy, information, components, assemblies, systems, or processes that become independently available to participate in causal interactions beyond the internal organisation from which they arise.

Outcome: A resulting state or condition arising from the participation of one or more outputs in subsequent causal interactions.

11.2 Function

The concept of function is widely used throughout systems science but has been interpreted in several different ways (Cummins, 1975). In everyday language, it often refers simply to what something does. Within General Systems Theory, however, function is more usefully understood in relational terms.

The function of a component, process, or system is its contribution to the behaviour, persistence, or capabilities of a wider organisation of which it forms part. Function therefore describes the significance of an activity within a broader system rather than the activity considered in isolation.

This distinction may be illustrated by a simple example. The heart contracts rhythmically throughout life. These contractions describe its behaviour. Its function, however, is not merely to contract but to circulate blood in a manner that contributes to the continued functioning of the organism. Similarly, the function of a pump within an engineered system is not simply to rotate an impeller but to contribute to the movement of fluid required by the wider system.

Function may therefore exist at many levels of organisation. The function of a process may contribute to the behaviour of a subsystem, the function of a subsystem may contribute to the behaviour of a system, and the function of that system may contribute to the behaviour of a larger parent system. Function is consequently a recursive concept that may be applied consistently throughout hierarchical organisations.

Functions are identified through their relationship to outcomes. A component or process performs a function because its outputs contribute to outcomes that support, modify, or otherwise influence the wider organisation. The same activity may therefore perform different functions when considered within different systems or for different analytical purposes.

The identification of function is partly ontological and partly epistemic. The causal relationships through which a component contributes to a wider organisation exist independently of observation. However, recognising those contributions as functions depends upon the way in which the wider system is represented and upon the analytical purpose of the investigation. Function is therefore neither an intrinsic property of an isolated component nor merely an observer attribution. It arises from real causal relationships represented within a particular systems context.

The definition adopted in this framework is therefore:

Function: The contribution made by the activity of a component, process, or system to the behaviour, persistence, capabilities, or outcomes of a wider organisation.

Because function describes a contribution to a wider organisation, it is inherently relational. The causal relationship through which the activity of a component, process, or system contributes to the behaviour, persistence, capabilities, or outcomes of a wider organisation may be described as a functional relationship. The same component may therefore participate in different functional relationships when considered within different parent systems or for different analytical purposes.

Functional Relationship: A causal relationship by which the activity of a component, process, or system contributes to the behaviour, persistence, capabilities, or outcomes of a wider organisation.

Function and purpose are closely related but distinct concepts. Functions describe the contributions that components, processes, or systems make to wider organisations. Purposes explain why certain systems perform those functions. Not all systems possess purposes. Many physical systems simply participate in causal processes, while purposeful systems organise their behaviour towards represented future states. Where no intrinsic purpose exists, observers frequently attribute purposes in order to summarise or explain system behaviour. The following sections develop these distinctions.

11.3 Purpose

Unlike function, purpose is not simply a consequence of a system's contribution to a wider organisation. Purpose refers to a desired or intended outcome represented by an agent and used to guide behaviour (Rosenblueth, Wiener, & Bigelow, 1943). Purpose therefore introduces intention into systems analysis and should be distinguished from causal consequence alone.

Purpose is not a universal property of systems but an emergent capability exhibited by certain classes of system. Systems may exhibit functions without possessing purposes, whereas purposeful systems necessarily exhibit both functions and purposes. General Systems Theory must therefore distinguish between systems that simply participate in causal processes and those capable of representing desired future states and modifying their behaviour accordingly. Purpose should consequently be regarded as an emergent systems concept rather than a foundational one.

A purpose is directed towards a future state that has not yet been realised. Systems capable of representing desired future states may modify their behaviour in ways intended to bring about those states. Purpose therefore involves both representation and action. It is not merely what happens, but what an agent seeks to make happen. An intrinsic purpose is therefore not simply a represented future state. It is the internally generated objective that organises and guides the behaviour of a purposeful system through its representation of that future state.

Purpose may exist at different levels of organisation. Individuals formulate personal purposes, organisations establish collective purposes, and societies may pursue shared purposes through institutions and governance. In each case, purpose is associated with agents capable of representing desired future conditions and acting intentionally in an attempt to realise them.

Purpose should not be confused with outcome. Outcomes describe the states or conditions that actually result from causal interactions, whether intended or unintended. Purpose describes the state or condition that an agent intends to achieve. Intended purposes may therefore be realised, only partially realised, or not realised at all. Likewise, systems frequently produce outcomes that were never intended.

The representation of purpose depends upon the analytical perspective adopted. An investigator may identify different purposes for the same system according to the level of organisation or the agents under consideration. Distinguishing purpose from function and from outcome therefore assists in avoiding teleological explanations that attribute intention where none exists.

The definition adopted in this framework is therefore:

Purpose: A desired future state or outcome represented by an agent and used to guide behaviour intended to realise that state or outcome.

Throughout this paper, the unqualified term *purpose* refers to intrinsic purpose. Where purpose is assigned externally by an observer, the term *attributed purpose* is used explicitly.

11.4 Attributed Purpose

The identification of purpose depends not only upon the properties of the system under investigation but also upon the perspective adopted by the observer. In many situations, purposes are attributed to systems by observers as a means of explaining or summarising their behaviour. Such attributions may be useful, but they should not be confused with purposes actually represented by the systems themselves.

Purpose is frequently attributed to biological, ecological, engineered, and social systems. For example, it may be said that the purpose of the heart is to circulate blood, the purpose of roots is to absorb water, or the purpose of a river is to drain a catchment. Likewise, the purpose of a bridge is often said to be transporting traffic. Strictly speaking, however, transporting traffic is the bridge's function. The purpose belongs to the designers, owners, or users who employ the bridge to achieve particular objectives. Such statements are often convenient summaries of the functions performed by these systems or components. However, they should not be interpreted as implying that hearts, roots, rivers, or bridges themselves possess intentions or represent desired future states.

Purpose may also be attributed at different levels of organisation. An observer may describe the purpose of a business as generating profit, while another may emphasise serving customers or providing employment. Similarly, the purpose of a government may be described in terms of security, welfare, justice, or economic development. Such differences frequently reflect differences in analytical purpose, perspective, or the level of organisation under consideration rather than contradictions about the underlying systems themselves.

Recognising attributed purpose as a representational concept helps to avoid teleological explanations that assign intention where none exists. It also reinforces one of the central themes of this framework: representations are constructed for particular analytical purposes and should be evaluated according to whether they faithfully represent the causal organisation of the systems under investigation.

The definition adopted in this framework is therefore:

Attributed Purpose: A purpose assigned to a system by an observer as an explanatory or interpretative representation, irrespective of whether the system itself has that purpose.

11.5 Purposeful Systems

Most of the concepts developed in this chapter apply to systems in general, irrespective of whether they possess awareness, intention, or agency. Purpose, however, represents an emergent capability exhibited only by particular classes of system. Systems capable of representing desired future states and modifying their behaviour in an attempt to realise those states may be described as purposeful systems (Ackoff & Emery, 1972).

Purposeful systems differ from other systems not because they cease to operate according to causal processes, but because internally represented future states become additional causal influences upon present behaviour. Purposeful behaviour therefore remains entirely consistent with a causal account of reality while recognising the emergence of new forms of causal organisation.

Purposeful systems may exist at different levels of organisation. Individual humans pursue personal purposes, while organisations may pursue collectively represented purposes through planning, decision-making, and coordinated action. Such systems remain subject to the general principles of systems theory but exhibit additional properties associated with representation, intention, learning, and adaptation.

The emergence of purposeful systems marks an important transition in the evolutionary development of complex systems. General Systems Theory provides the conceptual foundations for understanding such systems, while the detailed study of agency, motivation, governance, and collective action belongs to more specialised theories. The present framework therefore treats purposeful systems as an important class of system while recognising that the concepts required to understand them extend beyond the foundations established in this paper.

The definition adopted in this framework is therefore:

Purposeful System: A system capable of representing one or more desired future states or outcomes and modifying its behaviour in an attempt to realise them.

The emergence of purposeful systems therefore provides a natural transition from General Systems Theory to the study of social systems, where representation, agency, communication, governance, and productive coordination assume increasing importance.

12. Building Upon the Foundations

The preceding chapters have established the conceptual vocabulary and representational framework required to describe systems from a Critical Realist perspective. The concepts introduced have distinguished between ontology and representation, clarified the relationships between entities, systems, processes, states, behaviours, functions, and purposes, and established a coherent systems idiom applicable across multiple levels of organisation.

These foundations provide the basis for the next stage of the theory. The following papers build upon the concepts established here by examining the causal mechanisms through which systems interact, change, emerge, and evolve. The emphasis therefore shifts from defining the fundamental concepts of systems theory to explaining the causal processes responsible for the organisation and development of increasingly complex systems.

12.1 Systems as Causal Organisations

Throughout this paper, systems have been described as organised collections of interacting entities represented as wholes. From a causal perspective, however, systems may more fundamentally be understood as organised patterns of causal interaction. The components of a system do not merely coexist; they participate in recurrent causal interactions whose organisation gives rise to the behaviours, outputs and functions exhibited by the system (von Bertalanffy, 1968).

These causal interactions are realised through transfers of matter, energy, and information occurring between entities and processes. The organisation of these interactions gives rise to the behaviours, outputs, and functions exhibited by systems. Consequently, understanding systems ultimately requires an understanding of the causal relationships through which they are organised and maintained.

The present paper has established the conceptual foundations necessary to describe these causal organisations. The next paper will build upon these foundations by examining the nature of causal interactions, transfers, processes, and constraints in greater detail.

12.2 Constraints and Feedback

The behaviours exhibited by systems do not arise solely from their components or the causal interactions between them. They also depend upon the organisation that both enables and constrains those interactions and upon the conditions under which they occur.

Causal interactions may reinforce, regulate, or modify one another through the organisation of the system itself and through its interactions with the wider environment. Such influences are commonly represented as constraints (Ashby, 1956) and feedback (Wiener, 1948), concepts that play a central role in explaining the stability, adaptation, and evolution of systems.

The concepts developed in the present paper provide the conceptual foundations required to understand these ideas. The following paper will examine constraints, causal transfers, feedback, and related concepts in greater detail, showing how they regulate the behaviour of organised causal systems.

12.3 Emergent Properties, Behaviours and Outputs

The concepts developed throughout this paper provide the foundations for understanding emergence. As entities become organised into systems, the organisation of their causal interactions may give rise to new properties, new behaviours, and new outputs that are not exhibited by the individual components considered in isolation (Broad, 1925; von Bertalanffy, 1968).

Emergence is therefore not explained solely by the existence of components but by the organisation of the causal interactions between them and the conditions under which those interactions occur. The concepts of organisation, state, behaviour, outputs, function, and constraints established in this paper provide the conceptual vocabulary required to describe emergent phenomena consistently across different kinds of system.

The next paper will examine emergence in detail, distinguishing between emergent properties, emergent behaviours, and emergent outputs, and exploring the causal processes through which new forms of organisation arise.

12.4 Emergent Entities and Recursive Organisation

The emergence of new properties, behaviours, and outputs may, in turn, lead to the recognition of new entities that participate in further causal interactions. These emergent entities become components of higher levels of organisation, where they may themselves interact to produce new properties, new behaviours, and new outputs

(Simon, 1962; Koestler, 1967). In this way, systems become organised recursively into successive levels of causal organisation.

Recursive organisation is therefore a fundamental characteristic of complex systems. The same conceptual framework may be applied repeatedly across different levels of organisation, from physical structures to biological organisms, engineered systems, organisations, societies, and beyond. This recursive applicability is one of the principal strengths of General Systems Theory, providing a common conceptual language for describing increasingly complex forms of organisation.

The recursive emergence of progressively more complex causal organisations provides the basis for understanding the hierarchical organisation of reality. The following paper will develop this theme by examining the causal processes through which new forms of organisation emerge and become the foundations for further emergence.

13. Conclusion

This paper has sought to reconstruct the conceptual foundations of General Systems Theory from a Critical Realist perspective. Rather than introducing an entirely new systems theory, it has re-examined familiar concepts in order to clarify their meanings, distinguish their ontological and epistemic aspects, and organise them into a more coherent conceptual framework.

A recurring theme has been the distinction between reality and its representation. Throughout the paper, systems have been treated as real causal organisations existing independently of observation, while the concepts used to describe them have been recognised as purposeful epistemic representations developed for particular analytical purposes. This distinction has enabled concepts such as boundaries, openness, closure, state, behaviour, function, and purpose to be interpreted more consistently and with greater conceptual precision.

Throughout the paper, individual systems concepts have been developed not as isolated definitions but as successive stages in a coherent conceptual progression. The framework therefore provides both a systems vocabulary and an explanatory structure linking ontology, representation, organisation, and emergence.

Finally, the recursive nature of systems has been shown to extend beyond the organisation of reality to the concepts used to describe it. The same systems idiom may be applied repeatedly across successive levels of organisation, providing a common conceptual language capable of describing physical, biological, engineered, and social systems within a single coherent framework.

These foundations provide the basis for the subsequent development of the theory. Having established a coherent conceptual framework for describing systems, the emphasis can now shift towards explaining the causal processes through which systems interact, change, emerge, and evolve.

References

This paper is intended as a reconstruction of the conceptual foundations of General Systems Theory from a Critical Realist perspective rather than as a comprehensive review of the literature. References are therefore provided primarily to acknowledge the historical origins of major concepts where appropriate, rather than to document every aspect of the discussion. Many of the definitions, distinctions, and conceptual integrations presented here represent the author's own synthesis of ideas drawn from systems science, philosophy, engineering, and related disciplines.

Ackoff, R. L., & Emery, F. E. (1972). *On purposeful systems*. Aldine-Atherton.

Ashby, W. R. (1956). *An introduction to cybernetics*. Chapman & Hall.

Bhaskar, R. (1975). *A realist theory of science*. Leeds Books.

Bhaskar, R. (1978). *The possibility of naturalism: A philosophical critique of the contemporary human sciences*. Harvester Press.

Boulding, K. E. (1956). General systems theory—The skeleton of science. *Management Science, 2*(3), 197–208.

Broad, C. D. (1925). *The mind and its place in nature*. Kegan Paul, Trench, Trubner & Co.

Bruner, J. S. (1966). *Toward a theory of instruction*. Harvard University Press.

Cantor, G. (1895). *Beiträge zur Begründung der transfiniten Mengenlehre*. *Mathematische Annalen, 46*, 481–512.

Challoner, J. (2026). *Systems linguistics: Representation, compression, and productive co-ordination*. Rational Understanding. https://rational-understanding.com/EFGST/

Checkland, P. (1981). *Systems thinking, systems practice*. John Wiley & Sons.

Cummins, R. (1975). Functional analysis. *The Journal of Philosophy, 72*(20), 741–765.

Descartes, R. (1985). *Discourse on the method* (J. Cottingham, R. Stoothoff, & D. Murdoch, Trans.). In *The philosophical writings of Descartes* (Vol. 1). Cambridge University Press. (Original work published 1637)

Forrester, J. W. (1961). *Industrial dynamics*. MIT Press.

Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience, 11*(2), 127–138. https://doi.org/10.1038/nrn2787

Hall, A. D., & Fagen, R. E. (1956). Definition of system. *General Systems, 1*, 18–28.

Koestler, A. (1967). *The ghost in the machine*. Hutchinson.

Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society, 106*(6), 467–482.

Rosch, E. (1973). Natural categories. *Cognitive Psychology, 4*(3), 328–350.

Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Lawrence Erlbaum Associates.

Rosenblueth, A., Wiener, N., & Bigelow, J. (1943). Behavior, purpose and teleology. *Philosophy of Science, 10*(1), 18–24.

Saussure, F. de. (2011). *Course in general linguistics* (W. Baskin, Trans.). Columbia University Press. (Original work published 1916)

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27*(3), 379–423; *27*(4), 623–656.

Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society, 106*(6), 467–482.

Stevens, S. S. (1946). On the theory of scales of measurement. *Science, 103*(2684), 677–680.

von Bertalanffy, L. (1968). *General system theory: Foundations, development, applications*. George Braziller.

Wiener, N. (1948). *Cybernetics: Or control and communication in the animal and the machine*. MIT Press.

Definitions

D2.4 – Entity
Any identifiable portion of physical reality, including objects, fields, relationships, and events.

D2.24 – Characteristic (Property)

A feature common to entities used to classify or distinguish them.

D5.1 – Property Dimension

An epistemic representation of the possible positions that entities may occupy with respect to a selected property.

D2.25 – Variable

A property of an entity whose position on a property dimension may differ between entities or change through time.

D5.2 – Value

The symbolic representation of the position attributed to an entity on a property dimension at a specified time.

D2.26 – State
The set of properties and associated values that apply to an entity at a specified time.

D5.3 – Change

A difference in physical reality through time.

D2.27 – Change of State

A change in one or more of the properties or values that apply to an entity between specified times.

D2.28 – Continuous Change of State

A change of state in which a variable passes progressively through intermediate positions on a property dimension.

D5.4 – Discrete Change of State

A change of state in which an entity is represented as moving between distinguishable positions or states without intermediate positions being represented.

D2.19 – Relationship

A physically instantiated circumstance in which one entity may be described with reference to another through either their relative configuration or the transfer of matter, energy, or information between them.

D2.20 – Configurational Relationship

A relationship arising from the configuration (arrangement) of entities in space-time, without transfer.

D2.21 – Causal Relationship

A physically instantiated process involving the transfer of matter, energy, or information between entities.

D5.5 – Interaction

Participation by two or more entities in a causal relationship involving the transfer of matter, energy, or information between them.

D2.22 – Event

A temporally bounded causal interaction, or organised set of causal interactions, treated as an identifiable entity.

D5.6 – Process

An organised causal transformation in which inputs of matter, energy, or information are acted upon to produce outputs of matter, energy, or information.

D5.7 – Transfer

The physical passage or transmission of matter, energy, or information through a causal relationship.

D2.8 – Collection

A plurality of entities considered together on the basis of a shared feature, circumstance, or condition of selection.

D2.9 – Set

A collection represented as a single entity.

D2.12 – Set Duality

The principle that the same plurality of entities may be represented either as a collection of constituent entities or as a single set, depending upon analytical perspective.

D2.13 – Configuration

The overall spatial, temporal, or spatio-temporal arrangement of a plurality of entities.

D5.8 – Assembly

An identifiable entity comprising a plurality of causally interacting constituent entities.

D5.9 – Pre-Systemic Assembly

An assembly possessing organisation and structure but not exhibiting identifiable emergent properties, behaviours, or outputs.

D5.10 – System

An organised assembly exhibiting identifiable emergent properties, behaviours, or outputs and recognised and represented as a whole.

D5.11 – Organisation (Generic)

The network of relationships or interactions that brings entities together into a recognisable whole. This applies equally to aggregates, configurations, assemblies and systems.

D5.12 – Structure (Generic)

The overall pattern or arrangement exhibited by a network of relationships or interactions. Again, this applies throughout the hierarchy.

D5.13 – Organisation (Technical)

The network of causal interactions between the constituent components or processes of an assembly or system.

D2.14 – Structure (Technical)

The overall pattern exhibited by the causal organisation of an assembly or system.

D5.14 – Static Organisation

Organisation that changes little over the period of observation.

D5.15 – Dynamic Organisation

Organisation whose constituent entities, relationships, or interactions change through time while maintaining an identifiable overall organisation.

D2.15 – Static Structure

Structure that changes little over the period of observation.

D2.16 – Dynamic Structure

Structure that changes through time while maintaining an identifiable pattern.

D5.16 – System of Interest

The system selected and represented for investigation according to the purpose of the enquiry.

D5.17 – Component

An identifiable entity, process, assembly or system represented as contributing to the organisation, behaviour or outputs of a larger assembly or system.

D5.18 – Sub-system

A system that forms part of a larger system.

D5.19 – Super-system

A system containing one or more sub-systems.

D5.20 – Parent System

A system containing one or more sub-systems.

D5.21 – Emergence

The process by which the organisation and causal interaction of components gives rise to properties, behaviours, outputs, or causal capabilities that are not exhibited by the components considered in isolation.

D5.22 – Emergent Property

A property arising through emergence.

D5.23 – Emergent Behaviour

A behaviour arising through emergence.

D5.24 – Emergent Output

An output arising through emergence.

D5.25 – System Boundary

The epistemic representation that distinguishes a system of interest from its environment for the purposes of investigation.

D5.26 – Internal

Situated within the boundary of the selected system of interest.

D5.27 – External

Situated outside the boundary of the selected system of interest.

D5.28 – Environment

The set of external causal networks, including assemblies and systems, capable of causal interaction with a system of interest.

D5.29 – Component Environment

The external causal networks capable of interacting with a component.

D5.30 – Parent Environment

The environment of the component's parent system within which the parent system itself interacts.

D5.31 – Coupling

The degree to which the continued functioning of a subsystem depends upon causal inputs originating within a parent system rather than from other external causal networks.

D5.32 – Close or Tight Coupling

A high degree of dependence upon causal inputs generated within a parent system.

D5.33 – Loose Coupling

A relatively low degree of dependence upon causal inputs generated within a parent system, with greater dependence upon other external causal networks.

D5.34 – Input

Matter, energy, information, components, assemblies, systems, or processes that become available to participate in the causal organisation of a system.

D5.35 – Output

Matter, energy, information, components, assemblies, systems, or processes that become independently available to participate in causal interactions beyond the internal organisation from which they arise.

D5.36 – Internal Process

A process represented as occurring within the boundary of a selected system and participating in its causal organisation.

D5.37 – Input–Process–Output (IPO) Representation

A simplified representation of a system that describes its causal organisation in terms of inputs, internal processes, and outputs, providing a purposeful causal compression of the underlying organisation.

D5.38 – System State

A representation of the condition of a system at a specified time obtained from the values of selected system variables.

D5.39 – System State Change

A change in the representation of a system resulting from changes in the values of one or more selected system variables.

D5.40 – Behaviour

A representation of the characteristic patterns of state change exhibited by a system through time.

D5.41 – System Dynamics

A representation of the trajectories and longer-term patterns of system behaviour through time.

D5.42 – Stability

A representation of the tendency of a system to maintain or return to characteristic patterns of behaviour following disturbance.

D5.43 – Persistence

A representation of the continued existence of an identifiable causal organisation through time.

D5.44 – Open System

A system that participates in causal interactions with its environment through the exchange of matter, energy, information, components, assemblies, systems, or processes.

D5.45 – Closed System

A system in which one or more forms of causal interaction with the environment are absent or sufficiently restricted for the purposes of analysis.

D5.46 – Isolated System

A system that exchanges neither matter, energy nor information with its environment.

D5.47 – Experimental Closure

The deliberate restriction or control of causal interactions between a system and its environment in order to investigate selected causal mechanisms under defined conditions.

D5.48 – Analytical Closure

The deliberate restriction of the causal interactions represented within an explanation in order to simplify analysis while preserving those relationships considered relevant to the analytical purpose.

D5.49 – Relative Openness

The extent and character of the causal interactions between a system and its environment with respect to particular forms of exchange and a specified analytical purpose.

D5.50 – Outcome

A resulting state or condition arising from the participation of one or more outputs in subsequent causal interactions.

D5.51 – Function

The contribution made by the activity of a component, process, or system to the behaviour, persistence, capabilities, or outcomes of a wider organisation.

D5.52 – Functional Relationship

A causal relationship by which the activity of a component, process, or system contributes to the behaviour, persistence, capabilities, or outcomes of a wider organisation.

D5.53 – Intrinsic Purpose

A desired future state or outcome represented by an agent and used to guide behaviour intended to realise that state or outcome.

D5.54 – Attributed Purpose

A purpose assigned to a system by an observer as an explanatory or interpretative representation, irrespective of whether the system itself has that purpose.

D5.55 – Purposeful System

A system capable of representing one or more desired future states or outcomes and modifying its behaviour in an attempt to realise them.

D5.56 – Representational Role

A representational role is the manner in which an entity or process is represented within a particular investigation, such as the system of interest, a subsystem or a component. Representational roles depend upon the explanatory purpose of the investigation and do not alter the underlying causal organisation of reality.

Propositions

P5.1 – Ontology and Representation
Systems theory distinguishes between ontological causal organisations existing independently of observation and the epistemic representations used to describe and investigate them.

P5.2 – System–Process Equivalence
A system and a process may represent the same causal organisation viewed from different analytical perspectives. A system emphasises organisation, while a process emphasises causal interaction and transformation.

P5.3 – Ontological Primacy
Physical reality consists of entities and the relationships between them. Systems theory represents this organisation rather than creating it.

P5.4 – Progressive Organisation
Collections, configurations, assemblies, organisations and systems represent progressively greater degrees of organisation.

P5.5 – Progressive Representation
General Systems Theory represents progressively richer organisation through progressively richer conceptual representations.

P5.6 – Emergent Systems
A system emerges when organised causal interactions produce properties, behaviours or outputs that are not exhibited by its interacting components in isolation.

P5.7 – Emergent Causation
Emergence gives rise to new causal capabilities that are not available to the interacting components considered in isolation.

P5.8 – Recursive Organisation
Systems are recursively organised such that they may comprise interacting subsystems while simultaneously functioning as components of larger systems.

P5.9 – Analytical Perspective
The same portion of reality may be represented as a collection, configuration, assembly, component, subsystem, process or system according to the analytical purpose of the investigation and the causal organisation relevant to that investigation.

P5.10 – Boundary Relativity
System boundaries are analytical representations selected for a specified purpose rather than intrinsic features of reality.

P5.11 – Environmental Relativity
The environment of a system comprises the external causal networks capable of interacting with it and is represented relative to the boundary and analytical purpose of the investigation.

P5.12 – Relative Openness
Openness and closure depend upon the particular exchanges being considered rather than applying uniformly to all interactions.

P5.13 – Coupling Principle
The degree of coupling of a subsystem determines the relative importance of the external causal networks upon which its continued functioning depends and therefore influences the most appropriate representation of its environments.

P5.14 – Functional Emergence
Functions arise from organised causal interactions and therefore belong to systems rather than isolated components.

P5.15– Purposeful Behaviour
Purposeful behaviour occurs only in systems capable of representing desired future states and modifying their behaviour accordingly.

P5.16 – Representational Selection Principle

The representation of an entity within a systems investigation should be chosen according to the explanatory relevance of its emergent properties, behaviours or outputs to the system of interest.

P5.17 – Ontological Invariance Principle

Changing the representation of an entity from system to subsystem or component does not alter the underlying reality. It changes only the level of causal organisation explicitly represented within the investigation.

