Formalizing Heuristics: Cognitive Strategies for Decisions Under Constraint*

Timothy W. Coleman[1,2,*,†], John Bittner[1,2,†]

[1]SUNY University at Buffalo, NY, USA
[2]National Center for Ontological Research, NY, USA

Abstract

This article introduces the Heuristic Decision Ontology, a formal ontological framework for modeling heuristics as Directive Information Content Entities (DICE) within the context of medical triage. Grounded in Basic Formal Ontology (BFO 2020) and aligned with the Common Core Ontologies (CCO 2.0), the Heuristic Decision Ontology classifies cognitive shortcuts as actionable specifications that prescribe life-critical decision behavior. This paper presents the ontology's structure, semantic commitments, class hierarchy, and demonstrates its application through medical triage scenarios. We conclude with a discussion of the ontology's extensibility and its contribution to transparent and consistent modeling of heuristic-based medical judgment.

Keywords

Heuristics, Bounded Rationality, Medical Triage, Decision Theory, Cognitive Architecture, Applied Ontology, Basic Formal Ontology

1. Introduction

Imagine a medic in a mass-casualty scenario, racing against time to determine who needs care first and foremost. Using the Simple Triage and Rapid Treatment (START) triage method, they tag an unconscious but breathing patient as Red (needs urgent care), a vocal but stable patient as Yellow (can wait), and a walking patient as Green (minor injuries) [1]. These snap decisions rely on heuristics – quick, practical rules that cut through cognitive complexity by homing in on key signals. While the term "cognitive" is used throughout this paper, it refers to agent-neutral bounded reasoning strategies and does not imply human-exclusive implementation. From Aristotle's practical wisdom to modern triage rules, heuristics help experts act fast when time and data are scarce. While procedural models of heuristics exist, no prior work has provided a formally axiomatized, upper-ontology-aligned representation linking prescriptive content to decision processes in medical triage [2, 3].

Traditional clinical decision tools often stumble with messy, uncertain scenarios, relying on rigid logic or stats [4]. Our Heuristic Decision Ontology[1], built on Basic Formal Ontology (BFO) and the Common Core Ontologies (CCO) frameworks, models heuristics as reusable, machine-readable rules for triage and beyond, ready for smart systems [5, 6, 7, 8, 9].

Our ontology builds upon two foundational perspectives in cognitive science. The first is the heuristics-and-biases framework developed by Tversky and Kahneman, which highlights systematic biases that heuristics can produce under uncertainty [10]. To illustrate, triage personnel might exhibit an anchoring bias by sticking too closely to an initial impression of a patient's condition, or an availability bias by overestimating the severity of injuries that resemble recently encountered cases [10, 11].

The second is the ecological rationality framework advanced by Gigerenzer, which views heuristics as adaptive strategies applied to specific environments and not as sources of errors [12]. Taken from this vantage point, heuristics exploit the structure of the environment (e.g., the typical patterns of

Proceedings of the Joint Ontology Workshops (JOWO) - Episode XI: The Sicilian Summer under the Etna, co-located with the 15th International Conference on Formal Ontology in Information Systems (FOIS 2025), September 8-9, 2025, Catania, Italy,

*Corresponding author.
†These authors contributed equally.
timothywcoleman@gmail.com (T. W. Coleman); john.h.bittner@gmail.com (J. Bittner)
0009-0009-3931-4393 (T. W. Coleman); 00009-0001-5196-8279 (J. Bittner)

[1]Available at: https://github.com/TimothyWColeman/HeuristicOntology

injury severity in disasters) to yield adequate decision outcomes with limited resources. This aligns with Herbert Simon's notion of bounded rationality [13]. In a triage scenario, a heuristic like "treat the quietest patients first" (on the assumption that the most critical patients may be silent or unconscious) can be seen as ecologically rational: it works well given the environment of mass-casualty events, even if it deviates from a fully deliberative medical evaluation [14].

The Heuristic Decision Ontology bridges these two perspectives: one view sees heuristics as shortcuts that often lead to mistakes, and the other identifies efficient strategies to enable good decisions under constrained operating conditions – providing a firm foundation on which to ontologically represent heuristics as they appear in daily life.

2. Background

Cognitive heuristics have been studied extensively in behavioral economics, decision theory, and artificial intelligence (AI). Pioneering work by Daniel Kahneman and Amos Tversky in the 1970s introduced the heuristics and biases framework, which demonstrated that people often rely on mental shortcuts such as the availability heuristic or representativeness heuristic when making judgments under uncertainty [10]. These shortcuts can lead to systematic deviations from normative models of rationality; a phenomenon they termed cognitive bias. This work challenged the classical economic assumption of fully rational agents and catalyzed a paradigm shift in understanding human judgment. Kahneman later elaborated this perspective into the dual-system theory of cognition, distinguishing between System 1: fast, automatic, heuristic-driven processes and System 2: slow, deliberate, and rule-based reasoning [15].

While Tversky and Kahneman viewed heuristics as a sort of mental shortcut that assisted people in making decisions quickly, they focused more on how these shortcuts lead to predictable mistakes. In turn, they viewed heuristics as a simplification of judgment at the cost of accuracy. In contrast, Gerd Gigerenzer and colleagues advanced a complementary yet critical perspective. Gigerenzer argued that heuristics shouldn't be seen as shortcomings in human thinking, but rather strategies of efficiency that often lead to good outcomes. They should be perceived as tools that are applicable to specific environments. Through the ecological rationality framework, they argued that heuristics should not be viewed primarily as sources of error, but rather as adaptive strategies that work well in specific environments [12]. According to this view, heuristics exploit the structure of the environment to produce good-enough decisions with limited information and computational effort, which aligns with Herbert Simon's notion of bounded rationality [14, 16]. This approach suggests that the success of a heuristic depends not on its adherence to formal logic, but on its fit to the structure of the task environment.

These two perspectives are not entirely at odds with one another. Rather, it is important to see them as different sides of how heuristics work. In one, there is a focus on the risks associated with employing shortcuts. In the other view, attention is drawn to the practical value of those shortcuts in the real-world when information and time are constrained. Taken holistically, these perspectives help provide insight into how people manage complexity.

Our approach builds on this shared perspective by treating heuristics as formally representable entities that serve to prescribe behavior under operating constraints. The Heuristic Decision Ontology is a domain-specific ontology that captures common and well-defined cognitive approaches, aligning with McCaffrey and Wright's pluralist framework for cognition [17].

2.1. A Pluralist Approach to Prescriptive Heuristics

The Heuristic Decision Ontology aligns strongly with this pluralist vision by offering a domain-specific ontological structure for cognitive heuristics. In this way, it supports the pluralist agenda that McCaffrey and Wright promote, contributing a specialized yet formally rigorous characterization of what heuristics are. It transforms heuristics from loosely defined psychological tendencies into computationally actionable entities that can be queried, reused, and embedded in intelligent systems. The ontology characterizes existing and well-known heuristics while remaining extensible to incorporate and define new heuristics that emerge.

We align the Heuristic Decision Ontology to Basic Formal Ontology (BFO) 2020 [7], an ISO/IEC 21838-2 [18] standard, and the Common Core Ontologies (CCO 2.0) [19], to promote interoperability with other ontologies already in use in healthcare, defense, and AI domains [7]. BFO is a domain-independent upper-level ontology based on realist principles that distinguishes between continuants (entities that persist through time, e.g., a person or an instruction) and occurrents (events or processes that unfold over time, e.g., a triage assessment process) [7]. BFO provides a consistent scaffold for integrating diverse domain models and ensures that our representation of heuristics can co-exist with other ontologies. CCO is a widely used extension of BFO for representing domains of defense, intelligence, and manufacturing, among others.

2.2. Heuristics as Information Content Entities

Within this framework, the Heuristic Decision Ontology represents a heuristic as a Directive Information Content Entity (DICE), which is a specific kind of Information Content Entity (ICE). ICEs are classified as generically dependent continuants (GDC) in BFO, meaning they depend on some independent continuant to exist, but not on any specific one. For example, a GDC such as a heuristic instruction can be concretized in different media, such as a manual, a digital display, or a spoken utterance, without its identity depending on any particular bearer. While GDCs are often concretized in material entities, they can also be borne by other types of independent continuants, such as sites, depending on context and formal constraints [19]. An ICE is defined by the fact that it is ”about” something [19]. In our case, a heuristic (as a DICE) holds prescriptive content that lays the groundwork for an approach that an agent, such as a person or system, should carry out in a decision-making process. A DICE is specifically about a process and prescribes how that process should occur [19]. Examples of DICEs in general include control instructions, clinical guidelines, or protocols; any informational entity that prescribes actions rather than merely describing facts.

This paper uses the CCO 2.0 class DICE, which is defined as an ICE that prescribes some entity, such as an action or behavior [19]. While CCO builds on the general category of directive entities found in the Information Artifact Ontology (IAO), it introduces additional subclass distinctions that make it well-suited for modeling formal procedures, regulatory constraints, and operational rules. These distinctions support the aim of representing heuristics as structured, prescriptive content that can guide action in constrained decision-making contexts.

This raises an important distinction. While a best practice may also prescribe actions, it differs from a heuristic in scope and purpose. A best practice is typically based on evidence or consensus and is intended to work well under standard or optimal conditions. In contrast, a heuristic is used precisely when conditions are not optimal, when time, information, or cognitive bandwidth is limited. Heuristics are not necessarily the best choice in theory; they are the workable choice in practice under constraint. That difference is essential to how we model Heuristic Instruction as a DICE specifically designed to enable decision-making under pressure, not to maximize outcomes across all contexts.

Heuristics exist independently of any one person's mind [7]. For example, the rule ”If the patient is not breathing, open their airway; if they still do not breathe, tag Black (deceased)” is a heuristic rule in triage protocols. This rule exists as an informational entity. The medic may have internalized the rule into their mind by reading it in a triage manual, and it prescribes how they act in a specific situation. It is not just a habit in a medic's mind; it is an objective piece of content that can be communicated, stored, and analyzed. Thus, in the Heuristic Decision Ontology each specific heuristic (e.g., the Recognition Heuristic or Availability Heuristic) is identified as an instance of the class Heuristic Instruction, which is a subclass of DICE.

Additionally, it is important to note that we are not using the term ”heuristic” in its broadest computational sense. While heuristic algorithms like greedy search or hill climbing are common in computer science, our focus here is narrower. We are concerned with heuristics that act as prescriptive strategies: rules or procedures that guide agents in making decisions under conditions of constraint.

By ”cognitive,” we mean that the heuristic plays a role in structuring decision behavior, whether in a person or in an artificial system designed to operate under bounded rationality. Our intention is not to model every kind of heuristic used in optimization or search problems, but rather those that can be formally represented as DICE, specifically the kind that can be interpreted, applied, and evaluated within a process of decision-making.

Ontologically, a DICE is a kind of Information Content Entity that is about one or more processes. In the case of a heuristic, the content prescribes how a decision-making process should unfold under constraints. When an agent engages in such a process, their behavior may be guided by the DICE, even if the DICE is stored externally (e.g., in a manual) or internally (e.g., as a memory trace). It is important to note that while DICEs can guide or structure behavior, they are not realized in the BFO sense. Realization pertains to realizable entities like dispositions or roles, whereas directive content provides prescriptive structure without being instantiated through realization.

We can think of this decision-making process as one where an agent evaluates different options and picks a course of action. The directive content in the heuristic may include rules, criteria, or suggestions that shape and mold how an agent thinks and/or acts. It could also involve prioritizing certain information or evidence or using a heuristic approach when there is uncertainty.

In the Heuristic Decision Ontology, the heuristic's directive content is like a field manual or playbook for how an agent makes decisions. It lays the foundational rules, ideas, preferences, or biases that influence decisions. Specifically, this is captured as a Heuristic Instruction, a kind of DICE that shapes a Heuristic Decision-Making Process. It is designed to support an agent, for example, during a medical triage scenario when efficient choices are complicated, information is limited, and cognitive resources are constrained.

Under BFO's realism, heuristics unify biases (epistemic deviations) and adaptations (environmental fits) as DICE universals. Biases, such as availability leading to judgmental errors [10], and adaptations, such as recency fitting constrained environments [12], are instantiated as DICE particulars prescribing decision processes. This unification occurs because DICE, as universals, represent prescriptive content that can deviate epistemically (bias) or fit ecologically (adaptation) under bounded rationality [13]. Heuristic Instruction subClassOf DICE and prescribes some (Heuristic-Guided Process and (deviates_from some NormativeDecisionMakingProcess or exploits_structure_of some EnvironmentalStructure)), allowing instances to exhibit either deviation or adaptation depending on context [7, 19].

2.3. Procedural Models

Several existing cognitive architectures and efforts have addressed heuristics at a procedural or descriptive level. In this section, we briefly review their contributions and limitations in relation to our proposed ontological framework. Unlike previous approaches that implement heuristics, the Heuristic Decision Ontology treats them as semantically structured entities. Specifically, cognitive architectures, a computer-based model that mimics how people think and act, have sought to implement heuristics procedurally but rarely make them semantic entities, as seen in ACT-R [20] and Soar [21]. Both successfully operationalize heuristics as rules to simulate bounded rationality and task-specific strategies but lack structure for semantic reasoning or interoperability.

Poldrack and Yarkoni identify two core challenges in cognitive neuroscience: first, the difficulty of isolating well-defined cognitive functions, and second, the lack of structured mappings between mental constructs, behavior, and neural data [22]. To address this, they call for the development of formal cognitive ontologies that can clarify conceptual boundaries and enable integration across experimental paradigms. The Cognitive Atlas [23] is a taxonomy for classifying cognitive functions but lacks a formal top-level ontology like BFO, axiomatization, and support for reasoning. There is a need for rigorously axiomatized cognitive ontologies that clarify conceptual boundaries while enabling computational interoperability and formal inference.

3. Ontology Structure and Class Hierarchy

The Heuristic Decision Ontology was designed for semantic interoperability, axiomatic rigor, and machine reasoning over boundedly rational decision strategies. Development combined competency questions with reuse of existing ontologies The purpose was to ensure alignment with upper-level ontologies like BFO and CCO and support modeling heuristic decision strategies in computational terms.

To begin, several competency questions were formulated to define the scope of the ontology and to clarify what kinds of queries it should be able to answer. These included: *What kinds of heuristic decision strategies are relevant in time-constrained decision-making? How can we represent the relationship between a heuristic and the process that it informs? What distinguishes a heuristic-guided decision process from one based on intentional deliberation?* These types of questions served as guideposts throughout the design process to ensure that modeling decisions remained fully grounded in the ultimate use cases.

3.1. Types of Heuristic Instruction

At the center of the Heuristic Decision Ontology is the class Heuristic Instruction - a DICE that prescribes: a decision-making procedure intended to reduce cognitive effort under conditions of uncertainty or limited resources. Heuristic Instructions are prescriptive content about how one should decide under such conditions. For example, in the mass-casualty triage scenario introduced earlier, the medic's behavior followed such a decision strategy, resulting in the associated triage procedure.

The Heuristic Decision Ontology introduces four subclasses of Heuristic Instruction, representing families of heuristics characterizing distinct ways of guiding decision processes, as displayed in Figure 1 and described in Table 1.

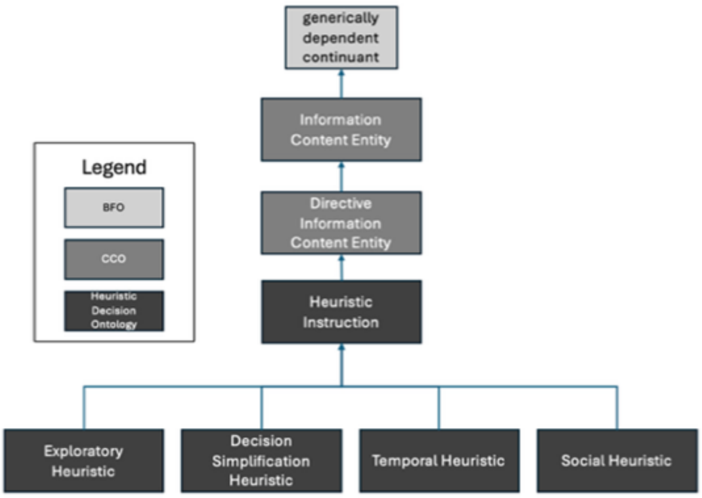

Figure 1: illustrates the structure of heuristic types.

The four major subclasses are: Decision Simplification Heuristic, Exploratory Heuristic, Social Heuristic, and Temporal Heuristic. Each of these contains specific heuristic types, such as Availability Heuristic (Decision Simplification), Curiosity Heuristic (Exploratory), Authority Heuristic (Social), Recency Heuristic (Temporal), etc.

A Decision Simplification Heuristic may assist a medic to home in on what is in front of them. For example, using the Availability Heuristic to flag a patient with chest pain as urgent because the medic remembers a recent case where a patient had similar symptoms and had a heart attack. This heuristic is all about leaning on what's easily recalled, making a call, while cutting through the noise to make a fast, efficient decision in chaos.

The Exploratory Heuristic nudges an agent to try something off the beaten path when the usual playbook of options doesn't fit. Returning to the medic in the mass-casualty scene, the Curiosity

Heuristic may push the medic to check an unconscious patient's odd symptoms like weird pupil reactions rather than just running through the standard vitals for triage evaluation. In this example, it is about chasing new clues when the situation is already murky, and the medic doesn't have a clear path forward.

A Social Heuristic comes into play when an agent is low on information and looking for direction from what others are doing. In the context of the medic, operating in a high-pressure mass-casualty event, the medic could notice that a seasoned colleague's approach seemed to be the go-to that all other medics were following. The Heuristic Social Proof may convince them to follow suit because if it is good enough for medics, it is probably solid enough to implement. It becomes a mental shortcut that takes stock in the wisdom of crowds and the group to make a decisive, reliable choice.

The Temporal Heuristic hinges on when information becomes available to the agent. For example, the Recency Heuristic might lead a medic to focus on the patients they just checked, especially one that suddenly looks worse, over others that they saw earlier. The freshest case feels like the most pressing, so this heuristic simplifies things by prioritizing what's most recent.

Table 1
Heuristic Instruction and Subclasses

Label	Definition
Heuristic Instruction	A Directive Information Content Entity that prescribes a procedure for decision-making characterized by the selective restriction of informational inputs, evaluative criteria, or processing operations, to enable execution of decisions under constraints such as limited time, incomplete data, or reduced capacity.
Decision Simplification Heuristic	A heuristic instruction that prescribes a procedure for reducing the complexity of a judgment or conclusion by focusing on readily available, representative, or pre-selected information, especially under conditions of time pressure, cognitive load, or incomplete data.
Exploratory Heuristic	A heuristic instruction that prescribes a procedure for selecting novel, unfamiliar, or varied options in order to seek information, stimulation, or potential reward, especially under conditions of habituation, monotony, or curiosity.
Social Heuristic	A heuristic instruction that prescribes a judgment or decision-making procedure in which an individual defers to, imitates, or aligns with the observable behavior, expressed choices, or inferred preferences of other individuals, particularly when personal information is sparse or uncertainty is high.
Temporal Heuristic	A heuristic instruction that prescribes basing a judgment or choice on the informational input processed most proximally to the moment of decision, under the assumption that this input is more relevant, accurate, or salient, particularly under conditions of limited memory, attentional constraint, or pressure to decide quickly.

3.2. Types of Heuristic-Guided Process

Complementing the representation of heuristics as informational entities, the ontology defines the class Heuristic-Guided Process, a BFO:process (an occurrent). This is a decision-making process that is structured, simplified, or otherwise directed by a corresponding Heuristic Instruction.

This class is crucial for modeling not only the content of heuristics but how they shape behavior. In other words, whenever an agent is using a heuristic to decide, that ongoing activity can be typed as a Heuristic-Guided Process. For example, a medic performing field triage rapidly under time pressure, if we recognize they are following heuristic rules, can be said to be engaging in a Heuristic-Guided Triage Process.

Heuristic-Guided Process serves as a bridge between the prescriptive nature of the heuristic (the content) and the dynamic, time-based nature of decision-making behavior (the process). These processes typically occur when an agent acts in accordance with a heuristic under conditions such as uncertainty, time pressure, or limited cognitive capacity, although heuristics may also be applied in routine, habitual, or training contexts where such constraints are not present. Each such process instance can be understood as the behavioral manifestation of a specific heuristic strategy, enacted in a constrained decision environment [14, 16, 12]. This framing aligns well with the concept of bounded rationality: the process is what bounded rational decision-making looks like in action, and the heuristic is the "rule" that bounds the rationality [16].

To capture the diversity of heuristic-influenced behaviors, we provide a range of named subclasses under Heuristic-Guided Process, each corresponding to a particular heuristic and thereby capturing a semantically distinct decision pattern. We further refine the classification by introducing an intermediate notion of a Heuristic Decision-Making Process, which is a Heuristic-Guided Process that specifically results in a decision outcome (as opposed to perhaps some other behavior). Under this, we categorize by the family of heuristics influencing it, such as Decision Simplification Process.

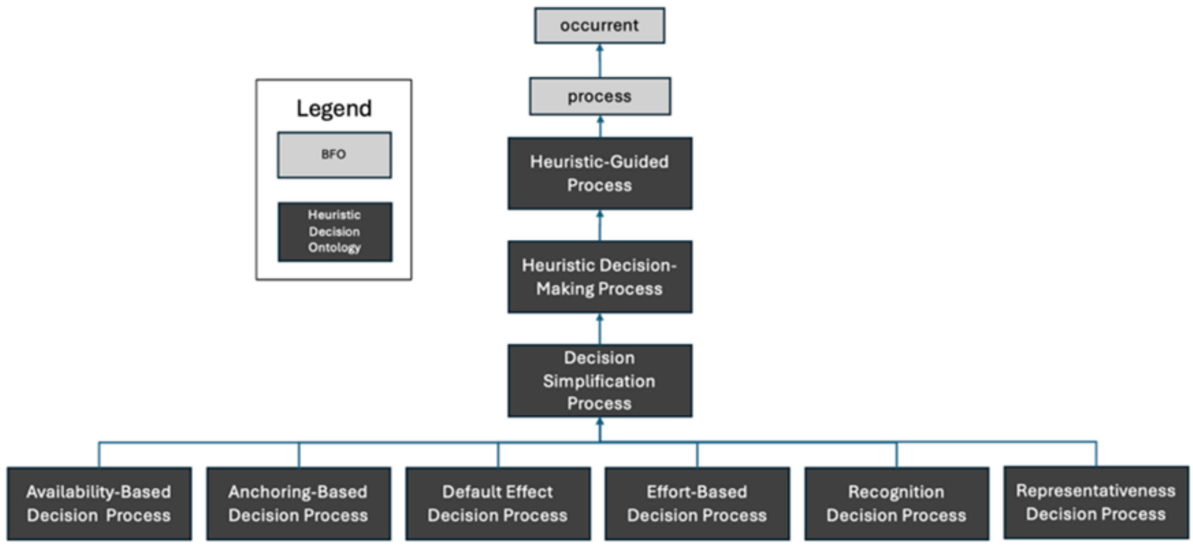

Figure 2: illustrates processes associated with heuristics decisions.

This hierarchical structure links the informational heuristics to behavior patterns in a formal, queryable way. Definitions incorporate relational 'prescribed by' to link processes to their directive heuristics. Table 2 illustrates core process classes.

Table 2: Heuristic-Based Decision Process and Subclasses

Label	**Definition**
Heuristic-Guided Process	A process prescribed by some heuristic instruction, characterized by simplified procedures for evaluation, selection, or action under conditions involving limited informational engagement, constrained time, or bounded capacity.
Heuristic Decision-Making Process	A heuristic-guided process prescribed by some heuristic instructions that culminates in a selection among alternative courses of action or outcomes, based on comparison criteria that are limited in scope, number, or complexity.

Continued on next page

Table 2 – *Continued from previous page*

Label	Definition
Decision Simplification Process	A heuristic decision-making process prescribed by some heuristic instruction that reduces the complexity of a judgment or conclusion by focusing on readily available, representative, or pre-selected information, especially under conditions of time pressure, cognitive load, and/or incomplete data.
Anchoring-Based Decision Process	A decision simplification process prescribed by some heuristic instructions that specifies or recommends the use of an initial reference value to structure subsequent evaluation or selection.
Availability-Based Decision Process	A decision simplification process prescribed by some heuristic instructions that prioritizes options that have been previously encountered over those not recognized or recalled.
Default Effect Decision Process	A decision simplification process prescribed by some heuristic instructions that results in the selection of a pre-specified option when no alternative is explicitly selected.
Effort-Based Decision Process	A decision simplification process prescribed by some heuristic instructions that restricts evaluation by prioritizing options requiring fewer observable or procedural steps to assess, acquire, or execute.
Recognition Decision Process	A decision simplification process prescribed by some heuristic instructions that prioritizes selection of a candidate encountered or accessed earlier in the process over candidates for which no such access event has occurred.
Representativeness Decision Process	A decision simplification process prescribed by some heuristic instructions that evaluates the similarity between a case and a category prototype, using that similarity to estimate probability or category membership.

A Heuristic-Guided Process is what happens when someone tackles a task with a streamlined approach because they're short on time, limited information, or mental bandwidth. A medic in a triage tent working through a flood of injured patients can't evaluate every detail, so they follow a quick, simplified method to decide to prioritize help first. While such strategies are often associated with human cognition, the ontology is designed to accommodate both human and artificial agents acting under comparable constraints. The ontology calls this a process shaped by Heuristic Instructions.

The Heuristic Decision-Making Process is a specific kind of Heuristic-Guided Process where the endgame is picking one option from a set of possibilities. It's prescribed by those Heuristic Instructions, which keep the decision focused by limiting what is considered. A parallel would be a medic deciding whether to send a patient to surgery or stabilize them on-site. They might use a heuristic to narrow down the choice, like focusing only on vital signs and injury type, making a call quickly under pressure.

A Decision Simplification Process takes this idea further by deliberately cutting down the noise. It's a Heuristic Decision-Making Process that's all about making complex choices manageable. For a medic triaging a victim, this might mean zeroing in on just a couple of key symptoms, like breathing trouble or heavy bleeding with the idea of keeping things simple.

An Anchoring-Based Decision Process is when that simplification hinges on a starting point that shapes everything else. The heuristic tells the agent or system to hold onto an initial value or idea and build a decision around it. Imagine a medic estimating how much fluid a dehydrated patient needs, and they might anchor on a standard amount they learned in training rather than calculating from scratch. That anchor keeps the decision quick and structured.

The Availability-Based Decision Process is focused on what's fresh or at the forefront of the mind. The heuristic pushes one to prioritize options you've seen before over unfamiliar ones. In triage, a medic might flag a patient with chest pain as critical because they just dealt with a heart attack case last week. The Availability Heuristic makes them lean on that recent memory, speeding up the call when constraints are imposed.

A Default Effect Decision Process happens when the heuristic nudges an agent or system towards a pre-set option if it doesn't actively choose something else. It's like a medic defaulting to a "Green" triage tag for a patient who seems stable unless clear red flags are evident. The Default Effect Heuristic essentially sets that baseline to keep decisions moving without overthinking every case.

The Effort-Based Decision Process is all about picking the path of least resistance. The heuristic tells an agent or system to choose options that take less work to figure out or to implement. For example, a medic might choose to stabilize a patient with a splint over arranging immediate surgery because it's quicker and simpler to execute, prescribed by an Effort Heuristic that favors a less intensive solution.

A Recognition Heuristic Decision Process is when an agent or system picks something familiar over the unknown. In a triage scenario, a medic might prioritize a patient with symptoms they recognize from past cases like a specific kind of burn over one with vague, unfamiliar complaints. The Recognition Heuristic makes that call faster by betting on what's known.

The Representativeness Decision Process is about judging something by how much it matches a typical example. The heuristic has an agent or system compare a case to a mental template to guess its category or likelihood. A medic might see a young patient with fever and rash and think "measles" because it fits the classic picture they know, prescribed by a Representativeness Heuristic. This process simplifies triage by slotting the patient into a familiar category without exhaustive tests.

By distinguishing between the informational structure of heuristics and the occurrent processes they influence, the ontology provides a modular and precise representation of boundedly rational decision-making. Though the focus is on medical triage, the pattern is general and can be applied across domains (e.g., interface design, autonomous vehicle decision-making, etc.), illustrating the ontology's cross-domain applicability. In the triage context, this structure allows us to formally represent scenarios like "a triage decision was made based on availability bias" as an individual of Availability-Based Decision Process – a subclass of Decision Simplification Process – linked to the Availability Heuristic. This opens the door for systems to perform audits (Was the decision process type appropriate?) and integration with clinical data.

Building on the earlier medical triage scenario where heuristics are demonstrated with a first responder, we can articulate a decision model using the Heuristic Decision Ontology. A medical responder encounters a patient experiencing chest pain and shortness of breath. The medic recalled a recent patient who deteriorated and had a heart attack. Based on this recollection (availability heuristic), the medic immediately classifies the current patient as high priority.

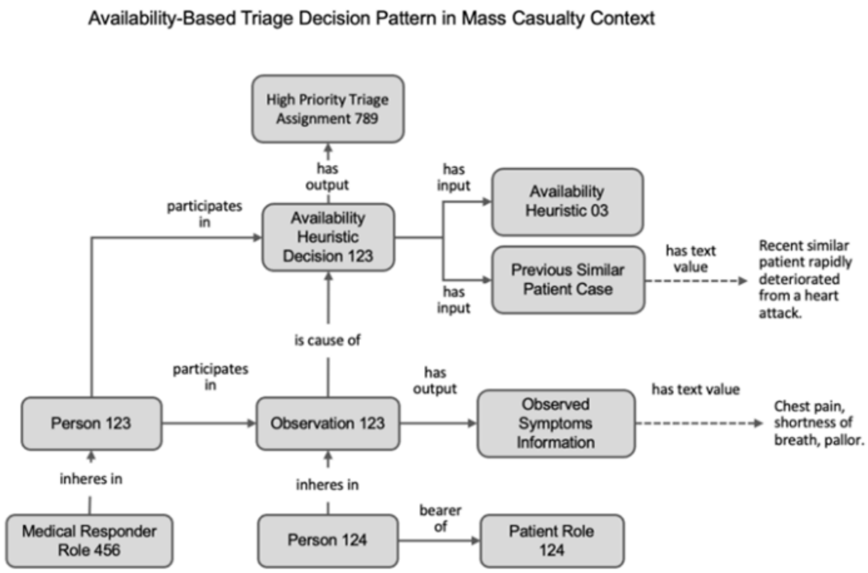

Figure 3: Example of an Availability-Based Decision Process modeled in the context of a mass-casualty triage scenario. The decision instance (AvailabilityHeuristicDecision123) is linked to the corresponding heuristic and patient presentation.

3.3. Exemplar Explanation

The specific decision (AvailabilityHeuristicDecision123) represents the medical triage event. The Availability Heuristic prescribed the decision, because the medic recalled a recent similar case of Person 124 and their rapid deterioration due to a heart attack. Details of the patient's clinical presentation are included ("chest pain, shortness of breath, and severe pallor"). The medic made the decision using the Availability Heuristic, assigning the patient as Immediate Priority (Red) due to concerns that they may rapidly deteriorate similarly.

These foundational modeling elements provide a coherent typology of heuristic strategies and their behavioral manifestations, enabling systems to perform basic classification, validation, and type-checking. For example, the Availability Heuristic is a subclass of Decision Simplification Heuristic, which is a subclass of Heuristic Instruction, and each corresponding behavioral manifestation is modeled as a subclass of Heuristic-Guided Process. Structuring the ontology this way helps keep different cognitive strategies distinct, reusable, and easy to apply across different systems.

The Heuristic Decision Ontology offers a formal ontological framework for a specific and underrepresented class of cognitive heuristics. In contrast to procedural or descriptive approaches like ACT-R and Soar, which simulate heuristics but do not define them semantically, the Heuristic Decision Ontology models heuristics as DICE that are both prescriptive and computationally actionable. This ontology helps clarify how decisions unfold and makes it easier to follow the logic behind them. It also improves compatibility across systems by giving them a shared way to evaluate decision-making. As a result, it brings abstracted decision theories closer to real-world use in AI and cognitive systems.

One way to improve the ontology further will be to add more detailed properties. These would help with classifying the heuristics further and more precisely to make it simpler to connect with other ontologies. For example, in a connected knowledge graph, one could show how an agent applies a specific heuristic during a process. This would allow for a clear, traceable link from the informational directive to the behavior it informs. On top of that, one could also add values for confidence level, response time, or cognitive load to support more detailed reasoning about how decision strategies work under specific operational constraints.

4. Properties and Reasoning Support

To enhance semantic expressivity and support richer modeling of heuristic-driven behavior, these properties enable queries such as identifying when a process both deviates from normative procedure and exploits environmental constraints. A set of object and data properties is proposed for future integration into the Heuristic Decision Ontology.

Table 3
Heuristic Decision Ontology Object Properties

Object Properties	Explanation	Domain	Range
prescribed by	Links a Heuristic-Guided Process to the Heuristic Instruction that directs its structure.	process	Heuristic Instruction
guides_behavior	Connects a Heuristic Instruction to the class of processes or actions it is designed to influence. This property avoids normative connotations while preserving directional semantics.	Heuristic Instruction	bfo: process
has_intended_outcome	Specifies the general type of result a heuristic aims to promote (e.g., faster decision, reduced cognitive load, satisficing choice).	Heuristic Instruction	cco: DICE
is_applied_to	Identifies the kind of entity that the heuristic is applied to, such as a person, option, system, or interface element.	Heuristic Instruction	bfo: continuant
deviates_from	Links a process to a normative decision process from which it departs, capturing bias as epistemic deviation (e.g., simplified inputs vs comprehensive norms).	bfo: process	NormativeDecisionProcess
exploits_structure_of	Links a process to an environmental situation whose constraints or patterns it leverages, capturing adaptation (e.g., time pressure in triage)	bfo: process	EnvironmentalStructure

These properties, consistent with CCO 2.0 patterns, support unification of biases and adaptations relationally, enabling inference like an Availability-Based Decision Process deviates_from a NormativeDecisionProcess but exploits_structure_of a mass-casualty EnvironmentalStructure. They are intended to enable semantic linkages between heuristic specifications, decision processes, contextual triggers, and behavioral outcomes that facilitate inference, validation, and simulation across intelligent systems and cognitive architectures.

Table 4
Heuristic Decision Ontology Data Properties

Data Properties	Explanation	Domain	Range
confidenceLevel	Represents the expected reliability or empirical effectiveness of a heuristic in a specific domain or task type.	Heuristic Instruction	Float [0.0-1.0]
expectedAccuracy	Provides a measure of how often the heuristic yields correct or effective outcomes in a target domain.	Heuristic Instruction	Float [0.0-1.0]
applicabilityScore	A context-specific scalar that captures how well the heuristic fits a given decision type or environment.	Heuristic Instruction	Float [0.0-1.0] or categorical rating 'high'

5. Heuristics in Cognitive Science and AI

The Heuristic Decision Ontology, designed to be compatible with BFO 2020 and CCO 2.0, represents heuristics as DICE linked by formal relations (e.g., prescribes, has input, has participant, and has output) to processes, situations, and outcomes. This structure, grounded in BFO's realist foundation, enables computational systems to perform semantic inference, validate decision strategies, and simulate heuristic behavior in agent-based models and cognitive architectures.

One important avenue for future development involves clarifying the boundaries between heuristics and other adjacent constructs, such as habits, protocols, or intuitions. While the ontology treats heuristics as directive, context-sensitive decision strategies, there may be cases where a frequently repeated heuristic becomes routine and indistinguishable from a habit. Similarly, context matters, and it plays a decisive role in whether a behavior qualifies as a heuristic: what is applied as a heuristic in one situation or in one domain (e.g., emergency triage) may not be appropriate or applicable in another (e.g., deliberative policy processes and planning). These ambiguities raise ontological challenges, as formal class definitions aim for clear categorical boundaries that real-world strategies do not always respect. Further empirical grounding, scenario-modeling, and community feedback will be critical in refining the typology of heuristics and resolving such gray areas. The ontology is designed with this in mind and is constructed with modular extensibility to accommodate precisely such iterative refinements over time.

The ontology's current scope focuses primarily on classification and representational clarity; it does not yet define Shapes Constraint Language (SHACL) constraints, closed-world validations, or domain-specific axiomatizations required for deployment in production-grade AI environments or high-assurance decision support systems. This paper also acknowledges that the typology of heuristics introduced, including Decision Simplification, Exploratory, Social, and Temporal – while grounded in cognitive science literature, is neither exhaustive nor fully canonical [10, 12]. Additional cognitive families may be needed to accommodate specialized domains such as military decision-making, consumer behavior, or clinical diagnostics.

6. Conclusion and Outlook

The Heuristic Decision Ontology offers a semantically structured and extensible framework for representing heuristic cognition in both human and artificial systems. Developed in alignment with BFO 2020 and structurally compatible with CCO 2.0, it models heuristics as DICE that prescribe cognitively simplified strategies under bounded conditions such as uncertainty, time pressure, and limited information. What distinguishes the Heuristic Decision Ontology is its ability to represent not only the structure of heuristics but also their behavioral manifestation through Heuristic-Guided Processes. This ontological separation between prescriptive information entities and occurrent decision behaviors allows for precise modeling of the interaction between cognitive shortcuts and real-world action. It supports bidirectional reasoning between heuristics and the processes they prescribe, enabling traceable, auditable, and ethically transparent modeling of decision logic across application domains. Ultimately, under BFO realism, the ontology unifies heuristics as DICE universals, reconciling biases as epistemic deviations [10] and adaptations as environmental fits [12], instantiated in prescriptive processes for constrained decisions via relational patterns that capture contextual deviation or exploitation without reifying qualities or dispositions. By formally linking prescriptive informational entities to their behavioral manifestations, this work advances cognitive ontology methodology beyond taxonomic enumeration to computationally actionable decision modeling.

Acknowledgments

The authors extend their gratitude to Barry Smith and John Beverley for critical feedback on early drafts of this work.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT to paraphrase and reword. After using this tool/service, the author(s) reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] U.S. Department of Health and Human Services, START adult triage, Webpage, n.d. URL: https://chemm.hhs.gov/startadult.htm, chemical Hazards Emergency Medical Management.
[2] Aristotle, Nicomachean ethics, Hackett Publishing, Indianapolis, 1999. Translated by Irwin, T.
[3] B. S. R. Pradelski, B. Tarbush, Satisficing equilibrium, arXiv, 2024. URL: https://arxiv.org/pdf/2409.00832. arXiv:2409.00832.
[4] T. J. Loftus, P. J. Tighe, A. C. Filiberto, P. A. Efron, P. Rashidi, S. C. Brakenridge, et al., Artificial intelligence and surgical decision-making, JAMA Surgery 155 (2020) 620–627. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7286802/.
[5] V. L. Patel, D. R. Kaufman, J. F. Arocha, Emerging paradigms of cognition in medical decision-making, Journal of Biomedical Informatics 35 (2002) 52–75. URL: https://doi.org/10.1016/S1532-0464(02)00009-6. doi:10.1016/S1532-0464(02)00009-6.
[6] J. N. Marewski, G. Gigerenzer, Heuristic decision making in medicine, Dialogues in Clinical Neuroscience 14 (2012) 77–89. URL: https://doi.org/10.31887/DCNS.2012.14.1/jmarewski. doi:10.31887/DCNS.2012.14.1/jmarewski.
[7] R. Arp, B. Smith, A. D. Spear, Building ontologies with Basic Formal Ontology, MIT Press, Cambridge, MA, 2015.
[8] C.-B. Benson, A. Sculley, A. Liebers, J. Beverley, My ontologist: Evaluating BFO-based AI for definition support, arXiv, 2024. URL: https://arxiv.org/pdf/2407.17657. arXiv:2407.17657.
[9] T. Coleman, J. Bittner, Heuristic Decision Ontology, GitHub repository, 2025. URL: https://github.com/TimothyWColeman/HeuristicOntology.
[10] A. Tversky, D. Kahneman, Judgment under uncertainty: heuristics and biases, Science 185 (1974) 1124–1131. URL: https://doi.org/10.1126/science.185.4157.1124. doi:10.1126/science.185.4157.1124.
[11] P. Croskerry, The importance of cognitive errors in diagnosis and strategies to minimize them, Academic Medicine 78 (2003) 775–780. URL: https://doi.org/10.1097/00001888-200308000-00003. doi:10.1097/00001888-200308000-00003.
[12] G. Gigerenzer, P. M. Todd, ABC Research Group, Simple heuristics that make us smart, Oxford University Press, New York, 1999.

[13] H. A. Simon, Rational choice and the structure of the environment, Psychological Review 63 (1956) 129–138. URL: https://doi.org/10.1037/h0042769. doi:10.1037/h0042769.

[14] G. Gigerenzer, W. Gaissmaier, Heuristic decision making, Annual Review of Psychology 62 (2011) 451–482. URL: https://doi.org/10.1146/annurev-psych-120709-145346. doi:10.1146/annurev-psych-120709-145346.

[15] D. Kahneman, Thinking, fast and slow, Farrar, Straus and Giroux, New York, 2011.

[16] H. A. Simon, Models of man: social and rational, Wiley, New York, 1957.

[17] J. McCaffrey, J. Wright, Neuroscience and cognitive ontology: A case for pluralism, in: F. De Brigard, W. Sinnott-Armstrong (Eds.), Neuroscience and philosophy, MIT Press, Cambridge, MA, 2022, pp. 427–466. URL: https://www.ncbi.nlm.nih.gov/books/NBK583725/.

[18] International Organization for Standardization, Information technology – top-level ontologies (TLO) – part 2: Basic Formal Ontology (BFO) (ISO/IEC 21838-2:2022), 2022. URL: https://www.iso.org/standard/74572.html.

[19] M. Jensen, G. D. Colle, S. Kindya, C. More, A. P. Cox, J. Beverley, The Common Core Ontologies, arXiv, 2024. URL: https://arxiv.org/pdf/2404.17758. arXiv:2404.17758.

[20] J. R. Anderson, C. Lebiere, The atomic components of thought, Lawrence Erlbaum Associates, Mahwah, NJ, 1998. URL: https://doi.org/10.4324/9781315805696. doi:10.4324/9781315805696.

[21] J. E. Laird, The Soar cognitive architecture, MIT Press, Cambridge, MA, 2012. URL: https://mitpress.mit.edu/9780262122962/.

[22] R. A. Poldrack, T. Yarkoni, From brain maps to cognitive ontologies: Informatics and the search for mental structure, Annual Review of Psychology 67 (2016) 587–612. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC4701616/pdf/nihms-684013.pdf.

[23] R. A. Poldrack, A. Kittur, D. Kalar, E. Miller, C. Seppa, Y. Gil, et al., The Cognitive Atlas: toward a knowledge foundation for cognitive neuroscience, Frontiers in Neuroinformatics 5 (2011) 17. URL: https://doi.org/10.3389/fninf.2011.00017. doi:10.3389/fninf.2011.00017.