

Ontology and Cognitive Outcomes

David LIMBAUGH^{a,1}, David KASMIER^a, Ronald RUDNICKI^b, James LLINAS^b,
and Barry SMITH^a

^a*Department of Philosophy, University at Buffalo*

^b*CUBRC, Inc.*

^c*Department of Mechanical and Aerospace Engineering, University at Buffalo*

Abstract. The intelligence community relies on human-machine-based analytic strategies that 1) access and integrate vast amounts of information from disparate sources, 2) continuously process this information, so that, 3) a maximally comprehensive understanding of world actors and their behaviors can be developed and updated. Herein we describe an approach to utilizing outcomes-based learning (OBL) to support these efforts that is based on an ontology of the cognitive processes performed by intelligence analysts.

Keywords. cognitive process, ontology, outcomes-based learning, intelligence analysis, machine learning

1. Introduction

We define ‘ontology’ as a controlled vocabulary of terms organized hierarchically, where the terms in the ontology are used to semantically enhance bodies of data in such a way as to make even highly heterogeneous data more consistently accessible to computers [1]. Herein we discuss the Cognitive Process Ontology and describe the role it can play in a system of outcomes-based assessment of analytical workflows, and how this system may then be applied to create a learning system to improve analytical outcomes. By ‘analytical workflow’ we understand a series of steps performed by humans or teams of humans using computers to transform heterogeneous data and information into decisions. The data in question can be in any domain, but we focus here on the intelligence domain, where workflows of the relevant sort have been intensively studied.

1.1. Background

As technology has advanced, the volume of data and information being made available to analysts has grown exponentially. Consequently, the Department of Defense (DoD) and the Intelligence Community (IC) must adapt to performing intelligence-related analysis on ever-growing amounts of distributed data [2]. By ‘intelligence-related analysis’ we mean more precisely the processing, exploitation, and dissemination (PED) portions of the IC’s *intelligence cycle*. This cycle begins with an *intelligence direction*, which is

¹Corresponding author, Department of Philosophy, University at Buffalo, 135 Park Hall, Buffalo, NY 14260; E-mail: dglimbau@buffalo.edu

established on the basis of a need for information. This is followed by *data gathering* from the operational environment. Gathered data is then *processed* and *exploited* for the sake of new information, and it is here that the analytic workflows that are the focus of this paper are located. Finally, new information is *disseminated* to those who need it, which results in the cycle beginning again [3]. This cycle is responsible for producing intelligence “at all levels of national security – from the war-fighter on the ground to the President in Washington” [3].

As the increasing use of technology increases the amount of data flowing through the intelligence cycle, this creates a situation in which technology becomes crucial also to the processing of these data, but this comes with challenges as a result of the fact that human analysts must also be involved [4,5]. Algorithms are intended to assist humans in handling ever larger amounts of data. However, the effectiveness of algorithms – even machine learning algorithms – is limited by how quickly important results can enter the human decision chain. This, then, requires that the outputs of algorithms be discoverable by, and comprehensible to, humans; this in turn requires that data formats, and data coding and tagging systems, are used that make human access to the data and human control over the data analysis process easier and the results of queries more readable. Furthermore, because an algorithm is merely a set of rules, it cannot create meaningful output data without meaningful input data. Not even the most sophisticated algorithm can produce meaning from a collection of, say, random numbers.

The human effort needed to address the above challenges, created by the sheer amount of data being collected, has widened the gap between data collection and data exploitation. The Director of National Intelligence has in consequence asserted:

Closing the gap between decisions and data collection is a top priority for the Intelligence Community (IC). The pace at which data are generated and collected is increasing exponentially – and the IC workforce available to analyze and interpret this all-source, cross-domain data is not. Leveraging artificial intelligence, automation, and augmentation technologies to amplify the effectiveness of our workforce will advance mission capability and enhance the IC’s ability to provide needed data interpretation to decision makers [6].

The upshot is that there is too much data and not enough human power to make effective use thereof.

1.2. A Proposal

The challenge is to close the gap between data collection and decision making. We propose to address this challenge by moving our attention one level higher, to the intelligence process itself. More specifically we propose

1. to collect data on the intelligence cycle as a whole and initially on the cognitive processes performed by humans and analogous processes performed by machines, and
2. to assess the effectiveness of different analytic workflows within the intelligence cycle described in terms of the cognitive processes involved.

A further goal is to use the collected data to automate portions of the intelligence cycle hitherto performed by human beings.

To collect, assess, and use data on the cognitive processes of humans in the intelligence cycle requires a means to organize data about such processes. To this end we have created the Cognitive Process Ontology (CPO), consisting of terms representing the cognitive processes – kinds of mental processes – used by analysts, such as ‘cognitive process of comparing’, ‘cognitive process of inferring’, ‘cognitive process of association’, and ‘analysis of competing hypotheses’ (ACH). Terms in CPO also represent the mental outputs of such processes, such as ‘Representation that is Believed’ and ‘Representation that is Warranted’.

Further ancillary terms are also included, representing what in the reality-outside-of-the-mind guides mental processes. These include terms, such as ‘indicator’, which refers to some portion of reality that, if known about, changes one’s estimation that some other portion of reality exists, has existed, or will exist. For example, knowing that a certain person visited a warehouse containing certain chemicals might be an indicator that an explosive device will soon exist.

Ancillary terms of this sort can be used to link CPO, which we can think of as an internally directed ontology, to other, externally directed ontologies such as those that comprise the Common Core Ontologies (CCO) ecosystem [7]. For example, we can use CPO-CCO combinations to create compound terms (and corresponding complex graphs) such as ‘Representation that is Warranted about Planned Missile Launch’ or ‘Evidence of Kinetic Kill Maneuver’. Aggregates of such terms can then be used to represent the components of investigative processes and process pipelines leading from data ingestion to informed decisions, including the outputs of such pipelines, for example in the form of predictions of real-world events [8].

Once we can represent investigative processes, including a granular representation of component mental processes, we are then able to tag such outputs of processes in light of how they contribute to decision making and fulfilling mission objectives. For instance, to what extent performing one cognitive process versus another cognitive process when analyzing satellite imagery affects the utility of the output of the analysis. Such tags, when used to tag sufficiently large amounts of process data, can serve as an aid in assessing which types and combinations of cognitive process are more likely to contribute positively or negatively to a decision-making process or mission outcome. Consequently, assessments will also suggest which, and how, cognitive processes should be managed so as to improve the intelligence cycle.

By annotating the investigative processes of analysts using machine readable ontologies, for example by utilizing logs of their computer operations, we can collect data both about how intelligence information came to be and also about what happens to that information in later stages of the intelligence pipeline. Over time this will lead to a more comprehensive understanding of the full intelligence-process workflow – what works, what doesn’t, why, and how to improve – drawing on the potential for outcomes-based research.

The goal of outcomes-based research is to find ways to promote those types of process workflows which have a higher likelihood of generating more useful outcomes. Research of this sort has demonstrated its value most conspicuously in information-driven biomedicine, where relative evaluations of treatment types are generated by associating data about the applications of such treatments to specific patients with data about subsequent outcomes [9]. To implement outcomes-based research in intelligence analysis, we need to make relative evaluations of intelligence processes and process workflows of

different types by associating data about instances of such processes with data about subsequent outcomes, and all of this with a focus on cognitive processes and how to stage them into the most effective overall analysis. This requires measuring the intelligence value generated by analytic workflows of given types, using their outcomes to determine average outcomes associated by workflows (cognitive process sequences) of those types, and then drawing conclusions from these average outcomes that allow assignment of metrics. A higher-level goal, once outcomes have been sufficiently measured, is determining which cognitive process workflow should be applied based on the features of an analysis task at hand, so as to achieve a desired outcome.

2. An Ontological Strategy

In building CPO we follow the methodology of ontological realism set forth in [10], according to which an ontology should be designed to represent entities in reality, including not only material things, their qualities and functions, and also the information artifacts used to describe and reason about them. As summarized in [11]:

This approach stems from the conviction that disparate ways of capturing data are best rendered interoperable by rendering them conformant to the ways things actually are. Realism implies, then, that any given assertion in an ontology can be evaluated on the basis of an objective criterion: Is the assertion true? this approach shifts ontology development away from the parochial concerns of particular implementations and toward expanded interoperability.

Every term in a realist ontology, accordingly, represents a type of entity that is instantiated by real-world instances of this type. The definition of this term captures what is common to all and only instances of this type. This definition has both a natural language form meaningful to human users and a logical form useful for machine processing. By semantically enhancing data with an ontology, both manual and automated methods can be applied to identify the relationships between entities of given types represented by given bodies of data and also to extrapolate new information based on those relationships, thereby complementing the human effort involved in delivering useful intelligence.

We are aware that many ontology-based approaches to data exploitation have failed. Our approach is based on a methodology developed in the field of bioinformatics, where ontologies – specifically the Gene Ontology (GO) and a series of ontologies built to interoperate with the Gene Ontology within the so-called OBO (Open Biological Ontologies) Foundry – are generally recognized as having been successfully applied [12]. This methodology has been piloted for DoD purposes by IARPA and the US Army Research, Development and Engineering Command, which sponsored a process of testing and validating by the International Standards Organization (ISO) and International Electrotechnical Commission (IEC) Joint Technical Committee Metadata Working Group. Two results of this piloting process are of relevance here:

1. International standard ISO/IEC 21838, approved in 2019 and scheduled to be published in 2020, including ISO/IEC 21838-2: Basic Formal Ontology (BFO), a top-level ontology to promote interoperability of domain ontology development initiatives [13], and

2. The Common Core Ontologies (CCO), a suite of interoperable ontologies based on BFO, including extension ontologies covering many defense and intelligence domains [7], are under initial consideration by the Mid-Level Ontology Ad-Hoc Committee of the InterNational Committee for Information Technology Standards (INCITS).

BFO is already being used as top-level architecture for a number of other ontology-building initiatives created under IC auspices, some of which involve use of the CCO, including work by the Applied Physics Laboratory (APL), the Institute for Defense Analysis (IDA) [14], and the Science Applications International Corporation (SAIC), as well as by some 300 ontologies development initiatives in medical, scientific and other areas [15].

A recent example of the utility of the realist-ontology approach for outcomes-based research can be seen in [16], which documents the building and implementation of the Ontology of Organizational Structures of Trauma systems and Trauma (OOSTT). By using OOSTT, data is able to be captured in real time and reasoned over by a machine, creating an ever-evolving representation of the domain. The information captured in this process is then able to be exploited to show, for example, correlations between the effectiveness of a trauma center and whether it is following regulations or has a particular organizational structure. Furthermore, the realist approach makes data organized by OOSTT interoperable with other data sets relevant to trauma centers, including data about patient consent [17] and patient outcomes [18].

2.1. An Ontology of Internal States

Almost all DoD- and IC-related efforts in ontology-building, for example [19], focus primarily on tagging data collected in external-world areas of interest such as geospatial locations, military units, sensors and their capabilities. Our goal here is to develop a complementary set of ontologies pointing internally, which is to say pointing to the thought processes (and analogous processes inside machines) involved in military and intelligence activities. For our present purposes we focus on the internal processes of single intelligence analysts. More specifically CPO focuses on providing the terminological resources for annotating data about mental processes of this sort which contribute to belief formation and belief changes (for example changes in confidence as to the veridicality of beliefs).

In building CPO, we reuse terms from existing ontologies wherever possible, but introduce new terms wherever needed. In either case, all terms are defined so as to be compliant with Basic Formal Ontology (BFO) as specified in [20]. Furthermore, because CPO draws its terms not only from CCO but also from other BFO compliant ontologies, there will be an effort in the interest of interoperability to bring those ontologies, whose content we want to re-use, into the fold of CCO extensions.

2.2. Analysis as a Feedback Loop

The collection and analysis processes we are addressing form feedback loops, as described in [21] and as illustrated in Figure 1. Feedback loops are iterative processes each additional iteration changing based on *feedback* from what happened in previous iterations. The feedback in this case is the intelligence gathered and processed in the prior

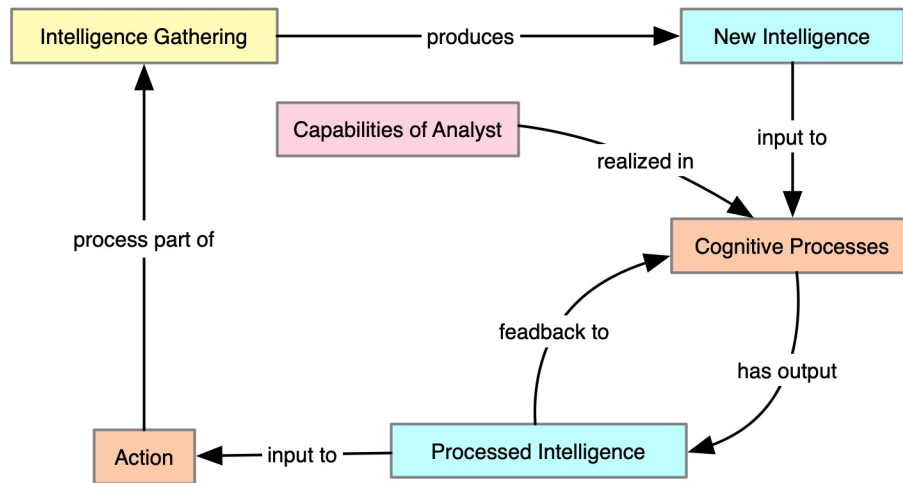


Figure 1. Intelligence Analyst Simple Feedback Loop

loop plus any still-relevant intelligence gathered in other previous loops. Gathered intelligence is processed within the loop by cognitive processes – the realizations of cognitive capabilities – possessed by analysts and that were largely acquired through training, study, and experience [22]. The outcome of processing intelligence includes new beliefs, new hypotheses, changes in confidence – all of which contribute to a growing body of processed intelligence – and which precipitate (are inputs to) actions like decisions relating to new intelligence collection steps, queries for further intelligence, or production and dissemination of results.

The proposal is that as the body of processed intelligence grows and various actions are taken, a traceable record of the steps taken within each iteration of a more complex loop will emerge. All paths actually taken through the loop will be recorded and related to outcomes to enable answering of questions like: “What sorts of queries prove useful in identifying better indicators?”, “What communication and documentation practices improve the timeliness of analytical output?”, and so on. As the traceable record grows in size it will become available also for more ambitious types of analysis, for example to support the creation of a catalog of lessons learned for use in training future analysts.

Importantly, outcomes include not only the degree to which an analytic workflow led to success or failure but also to what extent, and why, each component of the workflow contributed to this success or failure. For example, queries may be inadequate because they are poorly formed, addressed to the wrong recipient, duplicates, or already issued. Responses to queries, too, may be partial or inadequate. They may, provide only some of the information requested, address indicators incorrectly, provide only partial verification of an indicator, be inconsistent with background assumptions, or be inconsistent with previous responses.

3. The Cognitive Process Ontology

The Cognitive Process Ontology (CPO) is an extension of the Mental Functioning Ontology (MF) [23] and builds from the work on representations described in [24,25,26]. The central term of CPO – and the term that generally describes the work of an intelligence analyst – is ‘investigative process’, a subclass of what the Mental Functioning Ontology terms a ‘cognitive process’:

Cognitive Process =def. Mental Process that creates, modifies or has as participant some cognitive representation.

Investigative Process =def. Cognitive Process whose agent intends to establish or confirm that some portion of reality exists or does not exist.

An investigative process can be as simple as glancing to confirm the position of the hands of a clock and as complex as an extended International Criminal Police Organization (INTERPOL) terrorist hunt. Investigations unfold as an agent follows *indicators*, which are portions of reality (POR) that affect that agent’s estimation that some other portion of reality exists [27]. Practically anything (real) can be a portion of reality, and the same applies to those portions of reality that can serve, theoretically, as indicators. So not only are universals, formal relations, and instances portions of reality (and potential indicators), but so are combinations of these, such as a person of interest in Tucson, Italy having a meeting with a known arms dealer at 12pm GMT on October 12, 2014 [27].

Indicator =def. A portion of reality that, if it exists, affects our estimation that some other portion of reality exists.

3.1. Representations

Mental representations are what we *think with* when performing an investigation. Hence, what we are thinking about is determined by the content of our mental representations. To understand this class let us first introduce ‘Mental Quality’, which is a subclass of BFO ‘Quality’ [13].

Mental Quality =def. A quality which specifically depends on an anatomical structure in the cognitive system of an organism [24].

We are agnostic as to which parts of an organism constitute its cognitive system. We do however assume that it includes parts of the brain. The term ‘structure’ should also be understood in a very general sense, including for instance areas of the brain with particularly dense neuronal connections specialized to mental functioning of specific sorts.

The definitions of ‘system’ and ‘cognitive system’ presented here are provisional only, and should be read in conjunction with the proposed definition of ‘bodily system’ found in [28].

System =def. Material entity including as parts multiple objects that are causally integrated [29].

Cognitive System =def. System which realizes cognitive dispositions all of whose parts are also parts of a single organism.

As for mental qualities, we also assume that they are kinds of bodily qualities. But we remain agnostic as to what their physical basis might be, that is, what sort of independent continuant they inhere in. Mental qualities are either representational or they are not. Non-representational mental qualities include those that are responsible for giving emotional and sensational processes their characteristic feel. For example, the process of experiencing pain *hurts* because of the mental qualities involved in that process, and similarly for experiences of sorrow or joy.

3.1.1. *Mental Representations*

Mental Representations are Mental Qualities that concretize some Information Content Entity. In contrast to MF [23] and earlier versions of CPO [24,25,26] we do not here understand Mental Qualities as ever being about anything; this is to increase interoperability with suggested ways of modeling information with CCO [30]. Rather, the content of a Mental Quality is the Information Content Entity (ICE) it concretizes.

ICEs are BFO ‘Generically Dependent Continuants’ (GDCs), which means that an instance of an ICE can have multiple concretizations [13]. For example, the particular instance of an ICE that is *Structured Analytic Techniques for Intelligence Analysis* – also an instance of the subtype *textbook* – not only exists as concretized by the pattern of qualities inhering in the physical book (made of ink, glue, and paper) on your shelf, but also in the physical book on the shelf in the library, at the used bookstore, as well as concretized by the patterns of electricity that form the pdf file in your laptop. *Structured Analytic Techniques for Intelligence Analysis* is concretized in each case (they are all distinct copies of the same textbook). It is concretized by distinct instances of complex quality patterns inhering in different individual books or digital files. *Structured Analytic Techniques for Intelligence Analysis* thus depends generically on each and every book (or file) that concretizes it, and each and every book (or file) would have to be destroyed to successfully destroy *Structured Analytic Techniques for Intelligence Analysis* itself.

Mental qualities, in contrast to ICEs, are BFO Specifically Dependent Continuants: features of things that depend for their existence on some bearer. Thus, an instance of a Mental Quality specifically depends on part of a Cognitive System and is only located where its bearer is located, while the ICE concretized by a Mental Quality may also be concretized elsewhere.

When a quality concretizes some ICE, then that quality is a *Representation*, and when that Quality is also of the subtype Mental Quality, then it is a Mental Representation. We define each as follows:

Representation =def. Quality which concretizes some Information Content Entity.

Mental Representation =def. Representation which is a Mental Quality [24].

Mental Qualities are never strictly speaking about anything. However, Mental *Representations*, because of the content they concretize, are responsible for the intentionality (or directedness) found in a Cognitive Process. When asked, “What are you thinking about?” the answer is dependent on the Mental Representations had by the parts of the cognitive system that are participating in that process.

3.1.2. *Two Types of Concretization*

We distinguish two types of concretization: *original* and *derived* (compare Searle on ‘intrinsic’ and ‘derived’ intentionality in [31]). This distinction mirrors that between bona

fide and fiat boundaries; both types of boundaries exist and are genuine, but the former are associated with physical discontinuities, like walls and rivers, while the latter come into existence only through the intentional actions of agents, such as the signing of a legal document that specifies property lines [32]. Though the resultant property lines are products of fiat, they are nonetheless parts of reality and have legal significance.

3.1.3. *Original Concretizers*

Entities that are the original concretizers of some ICE are various types of Mental Representations. The concretization here is *original* because, like bona fide boundaries, the concretization relation between a mental representation and its ICE is in no way derived from the intentions of agents; that is, it is in no way derived from the way agents repeatedly use some symbol to communicate content; it is always the other way around.

3.1.4. *Derived Concretizers*

The paradigm entities that concretize derivatively are non-mental representations, as in symbols (quality-patterns) such as ‘dog’ or ‘ π ’, either spoken, written, or otherwise instantiated outside of the mind. A symbol begins to concretize some ICE because of how it is used (or intended to be used) by some intentional agent. Thus, it was only after an act of naming that the symbol ‘ π ’ became one way of expressing the ICE otherwise expressed as the symbol ‘pi’.

Following Chisholm’s doctrine of the primacy of the mental [33], the existence of a non-mental representation, which always concretizes its ICE derivatively, is explained in terms of the original concretization of that ICE by some Mental Representation and the intended use of symbols, by some agent, to increase the number of carriers of that ICE. (In what follows, if an expression is surrounded by single quotes (‘some expression’), then it denotes a *symbol*, and if surrounded by double quotes (“some expression”), then the *content* of some set of symbols is denoted instead.) The reason ‘ π ’ is associated with the ICE “the ratio of the circumference and the diameter of a circle” is because of, first, the original concretization of “the ration of the circumference and the diameter of a circle” in the mind of William Jones, who first introduced the symbol to carry the same ICE as ‘pi’, and later the original concretization in the minds of nearly every student who learned the language of mathematics. Before the intention of Jones for ‘ π ’ to be a vehicle for an ICE, ‘ π ’ was merely part of the Greek alphabet. Though original concretization always explains derived concretization, the temporal order can sometimes be backwards-looking. A photograph taken by a motion-sensor-activated camera concretizes an ICE about *the intruder*, not because we baptized it as such, but because we recognize it as such.

3.1.5. *Cognitive Representation*

Cognitive Representation is a subtype of Mental Representation (defined in 3.1.1 above). Because a Cognitive Representation always concretizes some Descriptive ICE, we can speak of a Cognitive Representation’s being correct, its degree of accuracy, and what it is about; though each of these is derivative on the representation’s concretized ICE. (A strategy for preserving interoperability between graphs that make the derivative nature of this relationship explicit and those that do not can be found here [34]).

A Descriptive ICE, then, is correct (henceforth; ‘veridical’) when it is about *the* portion of reality that it is intended to be about [24]. To say that a Cognitive Representation is veridical is to say that the represented POR – the POR the content of the Cognitive Representation is about – exists just as it is represented.

The distinguishing feature of a Cognitive Representation is what Searle called a ‘mind-to-world direction of fit’ [35,36]. A Cognitive Representations can be more or less accurate according to how well its content matches reality. If a Cognitive Representation is inaccurate, then the error is in the Cognitive Representation and not elsewhere; the Cognitive Representation aims to fit what it is intended to be about *in the world* and not vice versa.

Cognitive Representation =def. Mental Representation that has a mind-to-world direction of fit.

Contrast this with a type of Mental Representation that would be associated with a desire; a desire demands that the world fit it and not vice versa; it has a world-to-mind direction of fit.

3.1.6. *Representation that is Believed*

Some cognitive representations are taken by the agent to be veridical, whether they are actually veridical or not. Such a Cognitive Representation is what we referred to above with the term ‘Representation that is Believed’ (RTB). An RTB is treated by the agent (by her Cognitive System) as actually true, though it may or may not be actually true. More specifically, what distinguishes an RTB from a mere Cognitive Representation is that the latter is fused with a Positive Confidence Value (Compare what Meinong has to say about Ernstgeffhle or serious (or earnest) mental phenomena in [37].)

‘Fusion’ is a term adapted from Husserl [38] and is a primitive relationship that obtains between multiple Quality instances when they are so closely related that an additional Quality instance seems to emerge from them. Take for example what appears to be a solid green image displayed on a television screen, which upon very close inspection is actually colored by means of tiny yellow and blue squares, or pixels, thus giving a green appearance to the naked eye. The pixels are bearers of many instances of yellow and blue, and these instances appear to have fused into an additional instance of greenness. Similarly, when an instance of a Cognitive Representation and an instance of Positive Confidence Value are fused together in a cognitive system there seems to be an additional quality instance: an instance of an RTB.

A *Confidence Value* is a non-representational mental quality that, when fused with a Cognitive Representation, partially determines how that Cognitive Representation is utilized by a Cognitive System. The result is that the Cognitive System operates with that Cognitive Representation as if it is veridical. If Cognitive Representation CR2 that ‘My coffee is still too hot to drink’ is fused with a positive confidence value, then CR2 might be taken as input by the agent’s cognitive system when deciding to take a sip of the coffee. Because of CR2’s influence, the agent blows on the coffee first before taking a sip. Importantly, a fused Confidence Value should not be confused with second-order Cognitive Representations, such as a Representation about the likelihood of another Representation’s being veridical (as for example when you are asked: “Are you sure”) Such second-order mental representations are distinct from the pre-introspective

and non-representational confidence that we find fused with those cognitive representations which are RTBs.

Confidence Value =def. Mental Quality that, when fused with a Cognitive Representation CR, determines the extent to which a Cognitive System operates as if CR is veridical.

With this in mind we can now define ‘Representation that is Believed’ as follows:

Representation that is Believed (RTB) =def. Cognitive Representation that is fused with a positive Confidence Value.

(Note: a Confidence Value is represented as positive or negative in an RDF graph by means of data properties, like ‘CCO:has_decimal_value’, which are used to relate instances of types to numerical or nominal values. Accordingly, a Confidence Value would be positive when related by a data property to a decimal value of, say, above .5 in a Bayesian context or a nominal value of ‘high’, or even ‘positive’, in a nominal context.)

3.1.7. Representation that is Warranted

Following Plantinga [39], a Representation that is Warranted (RTW) is an RTB which holds an epistemically privileged place in a Cognitive System. It is so privileged because it was produced by some designed or vetted process so that, when in an environment of the sort that it was designed or vetted for, it reliably outputs veridical Cognitive Representations. As an analogy, consider an algorithm whose functioning is only designed and vetted for reliability for a certain type of input data. That algorithm would be functioning properly only when processing data of that type. Furthermore, the outputs of that algorithm should only be trusted when they are the product of the algorithm when it is functioning properly (or when the outputs are independently verified). As such, if an RTB is produced through a process of proper cognitive functioning, then it isn’t just *de facto* fused with a positive confidence value but also is such that it *should* be fused with a positive confidence.

Instances of such processes are instances of ‘Process of Proper Cognitive Functioning’ (PPCF), and the RTB’s formed by such processes are *warranted*:

Process of Proper Cognitive Functioning (PPCF) =def. Cognitive Process that has been successfully vetted or designed to reliably form veridical Cognitive Representations in environments of given types that include the environment in which the Cognitive Process is occurring.

Representation that is Warranted (RTW) =def. Representation that is Believed formed through Proper Cognitive Functioning in its vetted- or designed-for environment.

The privilege of an RTW is that it can justifiably be used in a Cognitive Process without further scrutiny. This is because, as shown in the graph in Figure 2, an instance of an RTW, by definition, is 1) produced by some PPCF and thus 2) rightly fused with a positive Confidence Value, which in this case is indicated by a decimal value representing a purported high chance (80%) of veridicality. Because these components are definitional, they can be enforced by reasoners, and thus it can be regulated that something only be tagged as an RTW if it has the requisite relations.

Furthermore, the introduction of warrant allows for a dimension of data integrity that goes beyond veridicality or confidence. For example, Figure 3 illustrates the formation of

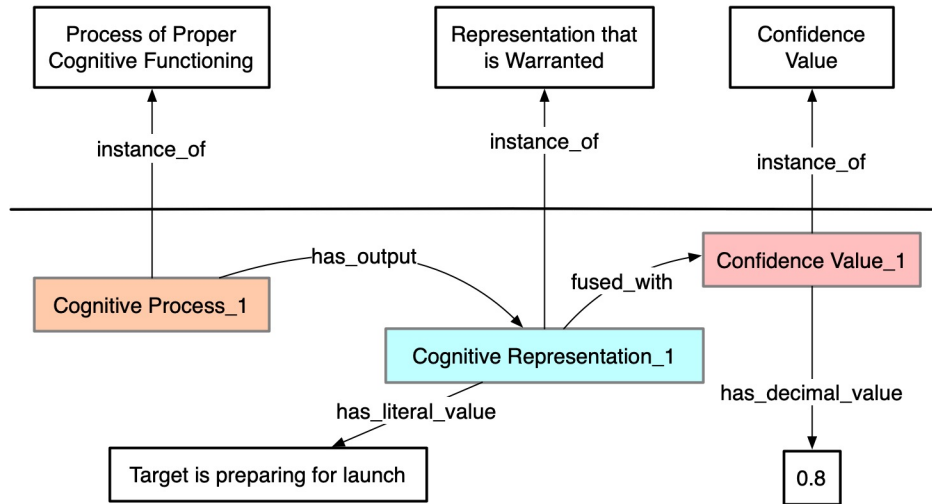


Figure 2. Above the horizontal line are terms from an ontology and below are items of data. A graph of a Representation that is Warranted including its entailments.

an RTB that may or may not be veridical but which the user holds with high confidence, although the Cognitive Process that outputted the RTB is not a PPCF. Here it is assumed that the relevant type of Cognitive Process requires veridical input data to be reliable; as such, even though the analyst may represent her RTB as fused with a high confidence value, the system knows to explicitly represent the information as unwarranted – as a mere guess (compare with [40]). The upshot is that warrant, veridicality, and confidence can each provide a dimension of data integrity to use when assessing information for the sake of decision making and outcomes-based research.

4. Implications for Future Research in Machine Learning and Automation

Here we lay out future research aimed at showing CPO’s capability to structure data and to deliver a method for discerning analysis processes that lead to significant knowledge gain from those that do not. This begins with synthetic data so as to prime the project for real data later, which is harder and more costly to acquire. These data should be about 1) a fictitious, but realistic, multi-domain mission, like a space and ground systems mission, 2) the intelligence collection processes that support that mission, and 3) the cognitive processes used by analysts to make progress via the collection processes. To avoid bias, all numeric data elements representing independent variables should be generated by using high quality pseudorandom number generators that are cryptographically secure, and all categorical data elements representing independent variables should be generated using state of the art randomization techniques such as those described in [41].

Next, subject matter experts determine the degree of correlation between events. While the measures of correlation should be created using a function from the independent variable to the dependent variable, a random value should also be added to the dependent variable using the same techniques that created the independent variable values.

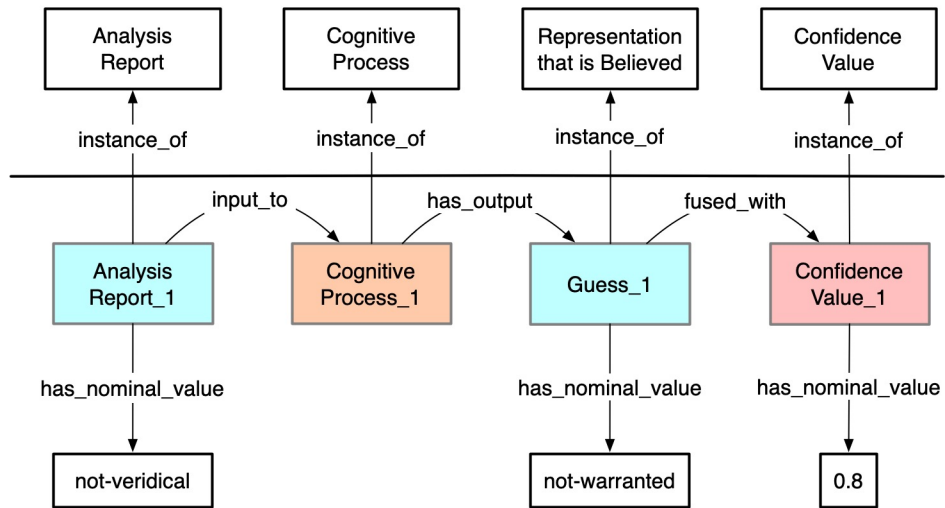


Figure 3. Above the horizontal line are terms from an ontology and below are items of data. A graph of a Representation that is Believed that is a mere guess because it was not formed through proper cognitive functioning due to the source information being not-veridical.

These formulas should not be known by the developers of the classifier and thus can reasonably be used as part of the assessment of the predictive model once produced.

Once the data is created and structured with CPO (and CCO), it can serve as the basis for the development of algorithms that classify intelligence analysis processes according to the amount of knowledge-gain they produce (that is, as training data for algorithm development). In particular a binary classifier should be used to assess whether or a not an intelligence analysis process results in a product that contributes knowledge gain to the mission. A binary classifier is a type of machine learning algorithm that produces a prediction model that, when given a new instance as input, places it into one of two categories. Such algorithms are usually supervised, meaning that the data from which they produce their initial model is already manually labelled as belonging to one of the two categories of interest.

5. Conclusion

We described methods for developing automation and outcomes-based learning in the intelligence domain. The use of ontologies provides a scalable means for retaining, integrating, and sharing a growing body of knowledge about the intelligence analyst's processes, including both immediate successes and failures and long-term outcomes. Moreover, these methods can be applied not only to cognitive data but also to data about analogous machine processes, like real-time sensor data streams, for example to identify data that would indicate an impending terrorist attack.

References

- [1] Salmen D, Malyuta T, Hansen A, Cronen S, Smith B. [Integration of Intelligence Data through Semantic Enhancement](#). In: Proceedings of the Conference on Semantic Technology in Intelligence, Defense and Security (STIDS 2011). Fairfax, VA: ceur-ws.org; 2011. 1
- [2] Lawrence A. K. [Sensor and Data Fusion: A Tool for Information Assessment and Decision Making](#). 2nd ed. Bellingham, WA: SPIE Press; 2012. 1
- [3] Rosenbach E, Peritz A. [Confrontation or Collaboration? Congress and the Intelligence Community](#). Harvard University: Belfer Center for Science and International Affairs. Cambridge, Massachusetts; 2009. 2
- [4] Marcus G. [Deep Learning: A Critical Appraisal](#). arXiv. 2018, arXiv:1801.00631. 2
- [5] Landgrebe J, Smith B. [Making AI meaningful again](#). arXiv. 2019, arXiv:1901.02918. 2
- [6] Office of the Director of National Intelligence. [The AIM Initiative: A Strategy for Augmenting Intelligence using Machines](#). Reports Publ. 2019. 2
- [7] CUBRC Inc. [The Common Core Ontologies](#). github.com. 2019 Available from: <https://github.com/CommonCoreOntology/CommonCoreOntologies> [cited 2020 Apr 10]. 3, 5
- [8] Chapman JR, Kasmier D, Limbaugh D, Gagnon SR, Crassidis JL, Llinas J, et al. [Conceptual Space Modeling for Space Event Characterization](#). Manuscript. Buffalo; 2020. 3
- [9] Clancy CM, Eisenberg JM. [Outcomes Research: Measuring the End Results of Health Care](#). Science. 1998;282(5387):245-246. 3
- [10] Smith B, Ceusters W. [Ontological Realism](#). Appl Ontol. 2010;5(3-4):139-88. 4
- [11] CUBRC Inc. [An Overview of the Common Core Ontologies](#). github.com. 2019 Available from: <https://github.com/CommonCoreOntology/CommonCoreOntologies/tree/master/documentation> [cited 2020 Apr 10]. 4
- [12] Kamdar MR, Tudorache T, Musen MA. [A Systematic Analysis of Term Reuse and Term Overlap across Biomedical Ontologies](#). Semant Web. 2017;8(6):853-71. 4
- [13] ISO/IEC DIS 21838-2 Information technology – Top-level ontologies (TLO) – Part 2: Basic Formal Ontology (BFO). International Organization for Standardization. Geneva, Switzerland; 2020. 4, 7, 8
- [14] Chan S, Haugh BA, Loaiza-Lemos FL, Wartik SP. [Department of Defense Public Safety and Emergency Management Communications: Interoperability Data Requirements](#). Alexandria, Virginia; 2017. 5
- [15] Smith B. [Users](#). basic-formal-ontology.org. 2018 Available from: <http://basic-formal-ontology.org/users.html> [cited 2019 Aug 9]. 5
- [16] Utech J, Judkins J, Otte JN, Colvin T, Rogers N, Rose R, et al. [OOSTT: a Resource for Analyzing the Organizational Structures of Trauma Centers and Trauma Systems](#). In: Proceedings of the Joint International Conference on Biological Ontology and BioCreative (ICBO-BioCreative 2016). ceur-ws.org; 2016. 5
- [17] Lin Y, Harris MR, Manion FJ, Eisenhauer E, Zhao B, Shi W, et al. [Development of a BFO-Based Informed Consent Ontology \(ICO\)](#). In: Proceedings of the 5th International Conference on Biomedical Ontology (ICBO 2014). Houston, TX: ceur-ws.org; 2014. 5
- [18] Ceusters W, Smith B. [Strategies for Referent Tracking in Electronic Health Records](#). J Biomed Inform. 2006;39(3):362-78. 5
- [19] Smith B, Malyuta T, Rudnicki R, Mandrick W, Salmen D, Morosoff P, et al. [IAO-Intel: An Ontology of Information Artifacts in the Intelligence Domain](#). In: Proceedings of the Eighth International Conference on Semantic Technologies for Intelligence, Defense, and Security (STIDS 2013). Fairfax, VA: ceur-ws.org; 2013. p. 33-40. 5
- [20] Arp R, Smith B, Spear AD. [Building Ontologies With Basic Formal Ontology](#). Cambridge, Massachusetts: The MIT Press; 2015. 5
- [21] Ford A. [Information Feedback and Causal Loop Diagrams](#). In: Modeling the Environment. 2nd ed. Washington, D.C.: Island Press; 2010. 5
- [22] Merrel E, Limbaugh DG, Smith B. [Mental Capabilities](#). In: Proceedings of the 10th International Conference on Biomedical Ontology (ICBO 2019). In: Proceedings of the 10th International Conference on Biomedical Ontology (ICBO 2019). Buffalo, NY: ceur-ws.org; 2019. 6
- [23] Hastings J, Ceusters W, Jensen M, Mulligan K, Smith B. [Representing Mental Functioning: Ontologies for Mental Health and Disease](#). In: Proceedings of the 3rd International Conference on Biomedical Ontology (ICBO 2012). Graz, Austria: ceur-ws.org; 2012. 7, 8

- [24] Smith B, Ceusters W. [Aboutness: Towards Foundations for the Information Artifact Ontology](#). In: Proceedings of the 6th International Conference on Biomedical Ontology (ICBO 2015). Lisbon, Portugal: ceur-ws.org; 2015. [7](#), [8](#), [10](#)
- [25] Limbaugh DG, Kasmier D, Ceusters W, Smith B. [Warranted Diagnosis](#). In: Proceedings of the 10th International Conference on Biomedical Ontology (ICBO 2019). Buffalo, NY: ceur-ws.org; 2019. [7](#), [8](#)
- [26] Kasmier D, Limbaugh DG, Smith B. [Foundation for a Realist Ontology of Cognitive Processes](#). In: Proceedings of the 10th International Conference on Biomedical Ontology (ICBO 2019). Buffalo, NY: ceur-ws.org; 2019. [7](#), [8](#)
- [27] Bittner T, Smith B. [A Theory of Granular Partitions](#). In: Applied Ontology. Frankfurt, GER: Ontos Verlag; 2008. p. 125-158. [7](#)
- [28] Smith B, Munn K, Papakin I. [Bodily systems and the spatial-functional structure of the human body](#). Stud Health Technol Inform. 2004;102:39-63. [7](#)
- [29] Mungall C. [RO:System](#) In: Relations Ontology. Ontobee. Available from: http://purl.obolibrary.org/obo/RO_0002577 [cited 2019 Apr 10]. [7](#)
- [30] CUBRC Inc. [Modeling Information with the Common Core Ontologies](#). github.com. 2019 Available from: github.com/CommonCoreOntology [cited 2020 Apr 10]. [8](#)
- [31] Searle JR. *Intentionality: An essay on the philosophy of mind*. Cambridge, UK: Cambridge University Press; 1983. [8](#)
- [32] Smith B. [Fiat Objects](#). Topoi. 2001;20(2):131-48. [9](#)
- [33] Chisholm RM. [The Primacy of the Intentional](#). Synthese. 1984;61(1):89-109. [9](#)
- [34] Dodds L, Davis I. [Qualified Relation](#). Linked Data Patterns. 2012 Available from: <https://patterns.dataincubator.org/book/qualified-relation.html> [cited 2020 Apr 10]. [9](#)
- [35] Searle JR. *Rationality in Action*. Cambridge, Mass: The MIT Press; 2001. [10](#)
- [36] Anscombe GEM. *Intention*. Ithaca, NY: Cornell University Press; 1963. [10](#)
- [37] Meinong A. *ber die stellung der gegenstandstheorie im system der wissenschaften*. Leipzig: R. Voigtlanders verlag; 1907. [10](#)
- [38] Husserl E. *Logical Investigatinos*. Translated by J. N. Findlay from the second German edition. New York: Humanities Press; 1970. [10](#)
- [39] Plantinga A. *Warrant and Proper Function*. New York, New York: Oxford University Press; 1993. [11](#)
- [40] Hogan WR, Ceusters W. [Diagnosis, misdiagnosis, lucky guess, hearsay, and more: An ontological analysis](#). J Biomed Semantics. 2016;7(1):1-15. [12](#)
- [41] Caiola G, Reiter JP. [Random Forests for Generating Partially Synthetic, Categorical Data](#). Trans Data Priv. 2010;3(1):27-42. [12](#)