Toward a General Theory of Knowledge

Luis M. AUGUSTO*

Editor-in-Chief

Journal of Knowledge Structures & Systems

November 2020

Vol.: 1 Issue: 1 Pages: 63-97

Abstract

For millennia, knowledge has eluded a precise definition. The industrialization of knowledge (IoK) and the associated proliferation of the so-called knowledge communities in the last few decades caused this state of affairs to deteriorate, namely by creating a trio composed of data, knowledge, and information (DIK) that is not unlike the aporia of the trinity in philosophy. This calls for a general theory of knowledge (ToK) that can work as a foundation for a science of knowledge (SoK) and additionally distinguishes knowledge from both data and information. In this paper, I attempt to sketch this generality via the establishing of both knowledge structures and knowledge systems that can then be adopted/adapted by the diverse communities for the respective knowledge technologies and practices. This is achieved by means of a formal–indeed, mathematical–approach to epistemological matters a.k.a. formal epistemology. The corresponding application focus is on knowledge systems implementable as computer programs.

Key words: Theory of knowledge (ToK); Science of knowledge (SoK); Industrialization of knowledge (IoK); Data, Information, & Knowledge (DIK); Knowledge structures; Knowledge systems

1 Introduction

In the last few decades, both society at large and individual persons or organizations became increasingly aware of the ubiquity of *knowledge* and access thereto. This awareness has caused the advent of a new age, the *Knowledge Age*, to be hailed as rapidly replacing the short-lived Information Age both in the more popular, business-oriented, circles and in the academe (e.g., Dzisah & Etzkowitz, 2012; Kidd, 2007; Ragsdell et al., 2002). In these, this phenomenon is particularly well reflected in the

*✉ luis.ml.augusto@gmail.com

recent proliferation of *"knowledge X"* labels, where X can be replaced by, among many other terms, *analysis*, *discovery*, *visualization*, *transfer*, etc. The interesting feature of this proliferation is not the terminological explosion itself–which is in fact a branching of the more primitive *Xs*, to wit, *engineering*, *representation*, and *management*–but rather the fact that each of these labels claims to capture a well-defined, separate discipline.

But short-lived as the Information Age might have been, it was solidly founded on a precise notion of *information*, namely C. E. Shannon's (1948); the hailed Knowledge Age so far lacks an equally solid foundation, with knowledge featuring in so many and diverse guises that experts from distinct *Xs*–henceforth *(knowledge) communities*–often talk past each other. Contributing to this state of affairs is the conception that defining *knowledge* is an application-dependent task and where applications are concerned we can do with merely "satisficing" notions of knowledge. For instance, J. Y. Halpern, a well-known scholar of knowledge representation, writes in the context of the debate over the right properties of knowledge that his "own feeling is that there is no unique right notion of knowledge; the appropriate notion is application dependent" (Halpern, 1995), and in a textbook on knowledge management one can read:

> A "good enough" or satisficing definition of knowledge has been shown to be effective ... once a satisfactory working or operational definition of knowledge management has been formulated, then knowledge management strategy can be confidently tackled. (Dalkir, 2005)

In contrast to this good-enough perspective, for centuries (actually: millennia), the word "knowledge," in particular its Greek or Latin equivalents, was pronounced with awe, being in certain periods even conceived as divine(-like): Plato famously spoke of the world of ideas (or forms), a "world"–the *huperuranus*–separated from this one we live in where the perfect and atemporal ideas of everything existing down here were to be contemplated by humans before their birth if they were to be able to know them, and Augustine, many centuries later but still highly influenced by this view, saw knowledge as an act of divine grace in which the "ideas," now placed in the *verbum*, were "communicated" to the knowing subject.[1]

Seen from a chronological viewpoint, between the erudite, often abstruse, discussions of knowledge in epistemology and the satisficing, good enough, notions allowed today there is the phenomenon that we can call, without discussing it further here, the *industrialization of knowledge* (IoK): It is all but impossible to find today a big company in which knowledge does not have a specific department, and everywhere there are staff employed as knowledge analysts, engineers, discoverers, architects, etc.

While this phenomenon can be seen as an assimilation of the concept of knowledge by the industry and the business world, the growing plethora of ever-diverging conceptions of knowledge has to be refrained. We need a notion of knowledge that is *invariant*, or *general* enough, in all these communities if they are to be able to communicate with each other. In particular, the conceptual promiscuity–indeed, often a confusion–of the concepts *data*, *information*, and *knowledge* (DIK), associated

[1] These are Plato's *theories of ideas* and *of (knowledge as) reminiscence* (cf. *Phaedo*, 80b, and *Phaedrus*, 245c-266c, respectively; see also the *Analogy of the cave* in *Republic*, 514a-520a), and Augustine's *theory of illumination*, ubiquitous in his *Confessions*. These are all highly complex texts originally written in Greek and Latin; if encouraged to pay them a visit with the help of translations into, and commentaries in, English, see Cooper (1997) and O'Donnell (1992).

with IoK, needs to be now effectively tackled. It is to this general theory of knowledge (ToK) that clearly distinguishes and defines knowledge that I call a *science of knowledge* (SoK; alternatively: *knowledge science*) proper.

In what follows, I firstly elaborate on why the current erudite notions of knowledge, albeit relevant, are not appropriate for the required generality and I begin to dissolve the DIK trinity, namely by distinguishing data and information. I then move on to the central task of this paper, to wit, the identification, definition from the viewpoint of generality, and segregation within DIK of knowledge structures and systems.

2 Beating about the Bush

IoK requires that we consider knowledge from three tightly connected factors: *Knowledge objects*, *knowledge agents*, and *knowledge processes*. The first fit generally into what I here call *knowledge structures*; together with the remaining two, these go on to compose what I call *knowledge systems*. Required is a notion of knowledge that is invariant with respect to all of these. This invariant notion needs also to be distinguished from the associated notions of *data* and *information*, often also confused. All these tasks are bound to conflict with representatives of the current state of affairs, and care is needed to avoid outright confrontation, reason why I see these tasks as "beating about the bush"; I will get to the point mainly in Sections 3 and 4.

2.1 Relevant Conceptions of Knowledge

Epistemology is the subject that concerns itself with the nature of knowledge and related topics. Unsurprisingly, given both the philosophical character and the longevity of this discipline, there are several opposing theories on what knowledge is–or is not. This diversity notwithstanding, some consensus has been reached by acknowledging the Platonic definition of knowledge as *justified true belief* as imposing the highest epistemic requirements, i.e. as containing both the necessary and the sufficient conditions for knowledge. In effect, a subject S can be said to know some proposition p if and only if (abbr.: iff) the following three conditions are satisfied:

(i) p is true.
(ii) S believes that p.
(iii) S is justified in believing that p.

While this *tripartite analysis of knowledge*, first formulated and discussed in Plato's *Theaetetus*, has moved the debate away from century-old metaphysical battlefields involving reality and the mind (see Augusto, 2005), it falls prey to a few challenges, in particular so to gettierization, the case that S may believe a true proposition p while not being justified in their belief and consequently fail to have knowledge with respect to p (Gettier, 1963). In effect, in mainstream epistemology the requirements for *epistemic justification* (condition [iii]) are often too high, satisfiable only by purely ideal subjects: For instance, S may be required to be able to reconstruct mentally the causal chain by means of which they found themselves in a doxastic relation with p (condition [ii]; see Goldman, 1967). This notwithstanding, the tripartite analysis is too pertinent to be summarily dismissed, as I argued for in Augusto (2011), and I shall retake it below, namely from the perspective of *formal epistemology*.[2]

A more recent field, *cognitive science*, also has its saying on what knowledge is. Although here, too, there is no single consensual definition of knowledge, this tends to be associated with *cognition* to the point that both terms are frequently interchanged (Augusto, 2010): *Cognition* is the processing (acquisition, storage, and retrieval) of information or data–ultimately: knowledge–from the environment with a view to securing the well-being of the cognitive, or knowledge, agent. While this definition more explicitly takes the agent into consideration, it is essentially applicable only to humans (and perhaps other animals) as knowledge agents. However, one of the innovating features of IoK is that it approaches cognition and knowledge also from the viewpoint of non-human agents, namely artificial agents, and I shall retake this topic below.

2.2 Dissolving the DIK Trinity. I: Data vs. Information

Faced with the trio composed of *data*, *information*, and *knowledge* (abbreviated: DIK), one often feels as perplex as medieval philosophy students when exposed to the notion of the trinity, composed of the three entities named the father, the son, and the holy spirit; like these, the elements that compose DIK are said to be one and the same and yet distinct. This distinction often assumes the form of a pyramidal hierarchy with data at the base, knowledge at the top, and information in the middle, making believe that somehow information intermediates between data and knowledge. But if perplexity can be a productive mental state in philosophy, in which field it is called *aporia,* it is simply undesirable in the context of IoK. In this, it allows for slogans such as "We turn your data into knowledge," meant to attract those, particularly companies, that possess large collections of data but are often at a loss about what to do with them. Luckily, we are here faced with an earthly trinity, one whose elements became confused in the development of the Information Age and the *aporia* can be dismissed. In this Section, where the task is mainly beating about the bush, information and data are compared and distinguished from each other; below, knowledge is tackled with respect to this distinguished duo.

Data – Let us suppose that we flip a coin several times and at each flip we write down the outcome: "H" if it is heads, "T" if it is tails. (Of course, we can also write "1" or "0," denoting heads or tails, respectively, for instance.) Given a sequence of 1, 2, ...k flips, we may end up with a collection of observations like the one in Table 1. Each of the pairs (1,H), (2,T), ..., abstracted as (F_i, X) for $i = 1, 2, ..., 9$ and $X = x \in$\{H,T\}, is a *datum*, and Table 1 is a collection of *data*.

Flip	1	2	3	4	5	6	7	8	9
Result	H	T	T	T	H	T	H	H	T

Table 1: Sequence of nine coin flips and respective outcomes.

Note the empirical nature of this collection: At each flip, we observed the outcome

[2]For both the many trees and the big forest that is epistemology, see, for example, Bernecker & Pritchard (2011). For a segregation of formal epistemology, see Hendricks (2006).

and noted it down in a convenient tabular form. Note also that before each flip F_i the outcome is a random variable X. This gives us the definition of a *datum* as a single value of a single random variable. A collection of data, denoted by Δ, is a *database*. A database can be of various forms, like Table 1, to be sure, but more often like the construct in Figure 1, called a *relational database*.

$$OUTCOME = \left\{ \begin{array}{c} (1, H), (2, T), (3, T), (4, T), \\ (5, H), (6, T), (7, H), (8, H), \\ (9, T) \end{array} \right\}$$

Figure 1: The database Δ_{Flip} for the coin-flipping sequence in Table 1.

Information – Suppose now that we face the outcome of flipping a coin from the viewpoint of uncertainty; for instance, we may win an amount of money if we get the outcome right, or lose money if we get it wrong. A sensible way to play this game is by thinking first in terms of the *probability mass function*, i.e. the function that gives the probability that a discrete random variable X is exactly equal to some value $x_i \in \mathcal{X}$ where $\mathcal{X}$ is an alphabet. Let $p : \mathbb{R} \longrightarrow [0, 1]$ be the probability mass function defined by

$$p_X(x_i) = \Pr(X = x_i)$$

where $-\infty < x_i < \infty$, Pr is a probability measure, $\sum p_X(x_i) = 1$, $p_X(x_i) > 0$, and $p_X(x) = 0$ for all other x. To make things simpler, let us consider the range of p to be a set of discrete values $\mathcal{X} = \{x_1, x_2, ..., x_n\}$. When flipping a coin, we can expect only two results, heads or tails. Thus, for X and $\mathcal{X} = (heads, tails)$, we have $\Pr(X = x_i) = 0.5$ for both $x_1 = heads$ and $x_2 = tails$. We have it then that the probability that flipping a coin will result in heads or tails is exactly 50% for each of the two possible outcomes. So far, there is in this datum no information proper: 50% is just a probability datum. Information proper enters the scene when we *quantify the uncertainty* involved in the value of the random value X as the outcome of a random process, a quantity that we can obtain by means of the formula

$$H(X) = -\sum_{i=1}^{n} p_X(x_i) \log_2 p_X(x_i)$$

called *entropy* (of the random variable X with probability mass function p). In the case of the random process of flipping a coin, we have

$$H(X) = -\left(\left(\frac{1}{2}\log_2\frac{1}{2}\right) + \left(\frac{1}{2}\log_2\frac{1}{2}\right)\right) = 1$$

and as $H(X)$ is measured in *bits*, the entropy of the random variable X in the random process of flipping a coin is exactly 1 bit. More specifically, $H(X)$ is the *self-information* of the random variable X; we can also measure the information (the reduction of uncertainty) due to another random variable Y, called *mutual information* and given by

$$I(X;Y) = H(X) - H(X|Y) = \sum_{x,y} p_{X,Y}(x,y) \log_2 \frac{p_{X,Y}(x,y)}{p_X(x)\, p_Y(y)},$$

where $P_{X,Y}(x,y) = \Pr(X = x, Y = y)$, a quantity that is always non-negative and symmetric in X and Y, i.e. X has as much information on Y as Y has on X, or $I(X;Y) = H(X) - H(X|Y) = H(Y) - H(Y|X) = I(Y;X)$.

Information vs. Data –The *bit* is thus the unit of information, just as the *datum* is the unit of data. These are rather distinct entities: The bit is always a quantitative measure, whereas the value of a random variable X in a datum can be nominal, quantitative, or qualitative (for instance, we could also write down the outcomes of flipping a coin as "good" or "bad"). In particular, the bit is a measure of both *data compression* and *data transmission*, the two main problems of communication theory, measured by $\min I\left(X;\hat{X}\right)$ (where $\hat{X}$ is an estimate of X) and $\max I(X;Y)$, called the data compression limit and the data transmission limit, respectively. We can thus say that a datum becomes information when it is *communicated*, i.e. compressed and/or transmitted or coded.[3] Inversely, we can say that information becomes data when there is decompression or decoding. However, what truly distinguishes both is entropy: Information is characterized by $H(X) > 0$ with respect to some random variable , whereas we speak of data when $H(X) = 0$. Equally, when there is no noise, i.e. when $I(X;Y) = 0$ for two random variables X and Y, we are in the context of data; of information whenever $I(X;Y) > 0$. Figure 2 shows a concise schema of the distinction data vs. information.

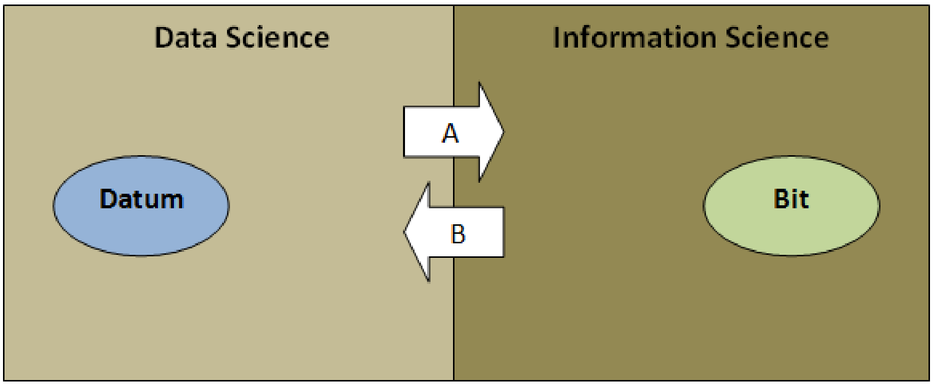

Figure 2: Data vs. Information. A: $H(X) > 0$; B: $H(X) = 0$.

Now, it will be useful to extrapolate to a less formal scenario. Let us suppose that someone, S, is gambling by means of coin flipping: If S correctly guesses the outcome of a flip (say, the first flip), S wins an amount of money; otherwise, S loses an equal amount of money. In other words, S is called to act upon the datum $(1, X)$. This datum is inert in itself; it is the need to act upon it, and consequently the need to quantify its uncertainty, that turns it into information. Because every action involves uncertainty, low as it might be, data become information when action is

[3]In effect, information theory was built upon the fundamental problem of communication, to wit, how to reproduce "at one point, either exactly or approximately, a message selected at another point" (Shannon, 1948). I am here assuming that information *theory* is at the core of information *science*. This apparently trivial remark is actually necessary, given the current lack of a consensual definition of information. See, for instance, Losee (2017).

called for upon them. In this particular example, S extracts from the datum $(1, X)$ the information that the value x for X is uncertain (by 1 bit), and thus S hopes luck to be on their side. Suppose now that S's guess was that $X = tails$. Then, the datum $(1, \mathrm{T})$ provides to S the information that S has won the amount of money gambled (but S could have lost it, and thus S acts happily), and the datum $(1, \mathrm{H})$, in contrary, provides to S the information that they lost their bet (but they could have won it, and thus are frustrated).

Importantly, note that nothing is being said about the *truth* or *falsity* of $(1, \mathrm{H})$ and $(1, \mathrm{T})$. The subject S is acting upon the information extracted from the *data* $(1, \mathrm{H})$ and $(1, \mathrm{T})$, regardless of their truth or falsity.

3 Knowledge Structures

In this Section, I elaborate on knowledge structures, from atomic structures and simple knowledge bases to compound structures and complex knowledge bases. These are considered in a purely static perspective, their dynamics being discussed in Section 4. I here retake the Platonic conception of knowledge as justified true belief, but now I shall get to the point. *Epistemic justification* is a very hard nut to crack, but mostly so in mainstream epistemology, where such bizarre entities as evil geniuses and clairvoyants can be invoked alongside with more earthly protagonists, such as dreams and the sensory modalities; in formal epistemology, it can be practically tackled via an adequate notion of *formal(ized) justification*. *Truth* can be an even harder nut to crack, but I shall take it as an element of formal semantics, where things are greatly simplified in comparison to mainstream epistemology. The adequate notions are here *interpretation* and *model*; in formal semantics, these are mathematical objects, and *truth* can be seen from an algebraic perspective. This remark is important, in order to lend generality to the elaboration below. I restrict this *formal(ized)* tripartite analysis of knowledge to beliefs that are assertions, or negations of assertions, proper; these are typically called *propositions*. The language $\mathfrak{L}$, a first-order language without function symbols and without the identity symbol (=), shall be considered an adequate formalism to this end. In this language, well-formed strings are generally called *formulas*; formulas without individual variables are called *sentences*, and propositions refer to either formulas with variables or sentences. I elaborate on this language in the Appendix. (For simplicity, I do not use the term "predicate.")

The formal language $\mathfrak{L}$ I shall be working with is a logical language because it has so-called *logical constants*, to wit, the elements of the sets Op and Qtf (see Appendix). However, one can work with non-logical operators and/or without quantifiers. As a matter of fact, one can work with a formal language that is as close as possible to a natural language.[4] This said, only a mathematical approach can give the generality I aim at, and I shall be working greatly from a viewpoint that is essentially that of mathematical logic.

[4]Formal semantics is typically associated with logic, but this is not necessarily so. In fact, it is first and foremost associated with formal languages, of which a logical language is a subtype (cf. Augusto, 2019). Actually, formal semantics also finds important applications in the study of natural languages. See Saba (2020; this issue) for a discussion of formal semantics as applied to natural languages in their relations to commonsense knowledge.

3.1 Facts and Knowledge Bases

3.1.1 Facts and models

Above, I considered the game of flipping a coin from the viewpoint of data and information. Let us now revisit the coin-flipping sequence above from yet another viewpoint, that of knowledge. Upon observation of the outcome of Flip 1, I pronounce or write down the following string of alphanumeric symbols:

(p) The outcome of Flip 1 is heads.

p is a *proposition*, and an *atomic* one at that. An atomic proposition just is a datum structured in a certain way, namely verbally in a natural language like English. Formally, an atomic proposition is often of the form "xRy", where R denotes a binary relation between any two arbitrary elements x and y. This can actually be written more compactly as the string $R(x,y)$, called an *atom* in formal lingo; hence, proposition p can be rewritten as the atom

$$(P) \qquad Outcome\,(1,h)$$

where $P \in \mathfrak{L}$, $\mathfrak{L}$ is a first-order logical language (see Appendix).[5] Clearly, a native English speaker will *understand* p effortlessly, and if trained in formal languages will also do so for P, but they might *not* understand

(p') El resultado del lanzamiento 1 es cara.

which in fact is the translation of p into Spanish and expresses just about the same, or has just about the same, *verbal meaning* as p.

Because verbal meaning depends to a great extent on natural languages and their associated cultures, it is highly variable, often varying between two speakers of the same language. Thus, some other kind of meaning is required if it must be invariant or general, allowing for a proposition p to be given meaning independently of the natural language in which it is formulated. What is needed here is an *interpretation* in the formal sense, to wit, a triple $\mathcal{I} = (\mathbf{U}, \Psi, \varpi)$, where $\mathbf{U} \neq \emptyset$ is the *universe of discourse* (e.g., Games), Ψ is a *signature function* mapping objects of the form $P = R(t_1, ..., t_k) \in \mathcal{P}_{\mathfrak{L}}$, where t denotes a *term* and $\mathcal{P}_{\mathfrak{L}}$ abbreviates $\mathcal{P} \subseteq \mathfrak{L}$, to a set of *distinguished elements* (e.g., $W_2 = \{0,1\}$), typically called *truth values*, and ϖ is a function assigning to every variable x an element in $\mathbf{U}$. It is frequently the case that $\mathbf{U}$ is subdivided into *domains* $\mathscr{D}_i \neq \emptyset$, $i = 1, ..., k$. In the case at hand, we have $\mathscr{D}_1 = \mathbb{N}$ and $\mathscr{D}_2 = \{h, t\} = \mathcal{X}$, where h abbreviates *heads* and t does so for *tails*. The domains are sets of *constants*, so that ϖ is an assignment of one or more constants to every variable. P is thus the result of, given the atom $Outcome\,(x,y)$, assigning to x an element of $\mathscr{D}_1$ and to y an element of $\mathscr{D}_2$:

$$Outcome \underbrace{(1,h)}_{\substack{\varpi(x) = 1 \in \mathscr{D}_1 \\ \varpi(y) = h \in \mathscr{D}_2}}$$

[5] A unary relation symbol $R(x)$ denotes a *property* (e.g., $IsWhite\,(x)$). *Outcome* actually abbreviates *IsOutcomeOf*. A relation must be written as a single string of one or more symbols. Henceforth, I write $p, q, ...$ to denote an (arbitrary) proposition and write $A, B, ..., P, Q, ...$ for emphasizing atoms. I shall denote an arbitrary formula by ϕ, ψ, or χ. (See Appendix.)

There is associated to $\mathcal{I}$ a *valuation function* $val : \mathcal{P}_{\mathfrak{L}} \longrightarrow W_n$ such that:

- $val_{\mathcal{I}}(x) = \varpi(x)$ for all $x \in Var_{\mathfrak{L}}$;
- $val_{\mathcal{I}}(R(t_1, ..., t_k)) = \Psi(R)(val_{\mathcal{I}}(t_1), ..., val_{\mathcal{I}}(t_k))$ for every k-ary relation symbol $R \in Rel_{\mathfrak{L}}$ and every symbol $t_i \in Ter_{\mathfrak{L}}$;
- $val_{\mathcal{I}}(o_i(\phi_1, ..., \phi_k)) = \widetilde{o_i}(val_{\mathcal{I}}(\phi_1), ..., val_{\mathcal{I}}(\phi_k))$ for every k-ary operator $o_i \in Op_{\mathfrak{L}}$ and every formula $\phi \in \mathcal{P}_{\mathfrak{L}}$ ($\widetilde{o_i}$ denotes the well-known logical construct called *truth table*, in turn an implementation of the function $\widetilde{o_i} : W_n^k \longrightarrow W_n$);
- $val_{\mathcal{I}}(q_j x\phi) = \widetilde{q_j}(distr_{\mathcal{I},x}(\phi))$ for every quantifier $q_j \in Qtf$, $j = 1, 2$, where $\widetilde{q_j} : \left(2^{W_n} - \emptyset\right) \longrightarrow W_n$ is a *truth*, or *distribution*, *function*, and $distr_{\mathcal{I},x}(\phi) = \left\{val_{\mathcal{I}_c^x}(\phi) \mid c \in \mathbf{U}\right\}$ is the *distribution* of ϕ in $\mathcal{I}$ with respect to x and $\mathcal{I}_c^x$ is an interpretation similar to $\mathcal{I}$ given $\varpi(x) = c$.

$\mathcal{I}_c^x$ corresponds to an *instantiation* of ϕ. For instance, let $\phi = \exists x Flies(x)$ and $\mathbf{U} =$ Birds; then, $\phi' = Flies(sparrow)$ is an instance of ϕ. Wrapping up what can be a long discussion (see Augusto, 2020d, Appendix), we have:

- "$A = \forall x R(x)$" is true iff, for all $c \in \mathbf{U}$, "$R(c)$" is true;
- "$A = \exists x R(x)$" is true iff, for some $c \in \mathbf{U}$, "$R(c)$" is true.

The interpretations considered in this Section are *Boolean*, i.e. they are interpretations over the truth-value set $W_2 = \{0, 1\}$, where 0 denotes *falsity* and 1 denotes *truth*. Furthermore, they are *classical*: A formula ϕ is either true or false *(principle of bivalence)*, and either ϕ or $\neg\phi$ is true (*principle of excluded middle*).[6]

Given both a set of formulas $\mathcal{P}_{\mathfrak{L}}$ and a valuation $val_{\mathcal{I}}$ for some interpretation $\mathcal{I}$, one can give meaning to all the formulas of $\mathcal{P}_{\mathfrak{L}}$ by means of a *model for* $\mathfrak{L}$, a pair $\mathcal{M}_{\mathfrak{L}} = (\mathcal{P}_{\mathfrak{L}}, val_{\mathcal{I}})$ such that for every formula $\phi \in \mathcal{P}_{\mathfrak{L}}$ it is the case that either $val_{\mathcal{I}}(\phi) = 1$ or $val_{\mathcal{I}}(\phi) = 0$. Whenever $val_{\mathcal{I}}(\phi) = 1$, it is said that there is a *model for* ϕ, a pair $\mathcal{M}_{\phi} = (\phi, 1)$ where $1 = val_{\mathcal{I}}(\phi)$.[7]

Suppose now that p is *true*, i.e. it is indeed the case that the outcome of Flip 1 is heads, or more formally, there is a model for $P \in \mathcal{P}_{\mathfrak{L}}$. Then, p is a *fact*, a unit of knowledge. A collection of facts corresponding to atomic propositions $\Xi = \{p_i\}_{i=1}^k$ is called a *(simple) knowledge base* (KB). Figure 3 shows the KB Ξ_{Flip} constituted by the facts corresponding to the data in Table 1.

Importantly, to say that p *is true* (or that P is *true in a model for* $\mathfrak{L}$) does not entail that this proposition is true in the sense that there is a certain relationship between it and reality that makes p *true* in the sense that this word is commonly used in the natural language English (e.g., "It's true; I actually saw it."). In effect,

[6] An interpretation $\mathcal{I}$ can be over W_n for $n \geq 3$, including $n \in [0, 1]$ (see Augusto, 2020c). Classical interpretations, however, offer many advantages, and any interpretation $\mathcal{I}$ over $W_{n \geq 3}$ can in fact be reduced to an interpretation over W_2 (see below).

[7] From a strictly algebraic viewpoint, a model is an algebraic structure, a pair (A, R) where $A \neq \emptyset$ and R is some relation over A. $\mathcal{M}_{\mathfrak{L}}$ and $\mathcal{M}_{\phi}$ are indeed models in this sense, as *val* defines a binary relation $\leq$ over elements of $\mathcal{P}_{\mathfrak{L}}$ such that for any two formulas $\phi, \psi \in \mathcal{P}_{\mathfrak{L}}$ we have $val(\phi) \leq val(\psi)$, namely via the chain $\mathbf{2} = \{0, 1\}$ where $0 < 1$. In particular, R is the *consequence relation* $\models$, a central construct of formal logic such that we write $\models_{\mathcal{M}} \phi$ (or $\mathcal{M} \models \phi$) whenever there is a model for ϕ. (See Augusto, 2020d.)

```
Outcome(1,h).
Outcome(2,t).
Outcome(3,t).
Outcome(4,t).
Outcome(5,h).
Outcome(6,t).
Outcome(7,h).
Outcome(8,h).
Outcome(9,t).
```

Figure 3: The knowledge base Ξ_{Flip}.

when one says that a proposition such as "Snow is white" is true in $\mathfrak{L}$, one is not referring directly to *snow* in the world, which is typically perceived as *being white*; one is simply referring to a construct of the formal language $\mathfrak{L}$ such that there is an interpretation $\mathcal{I}$ over which this construct (or its corresponding atomic formula $IsWhite\,(snow)$) is valuated as true. This requires that there be an *object language* of the type $\mathfrak{L}$, whose expressions are solely formulas, and a *metalanguage for* $\mathfrak{L}$ that, besides containing all its formulas, contains additionally expressions such as "is true" and "iff." Let us segregate the object language $\mathfrak{L}_0$ with respect to its metalanguage $\mathfrak{L}_1$, $\mathfrak{L}_0 \subset \mathfrak{L}_1$. Then, $\mathfrak{L}_0$ contains only *sentences* (e.g., $IsWhite\,(snow)$) while $\mathfrak{L}_1$ *additionally* contains *names* (e.g. "Snow is white"), more specifically a name for every sentence of $\mathfrak{L}_0$. Let us denote an arbitrary sentence by p and its name by "p". Then, an expression such as

$$\text{“Snow is white” is true iff } IsWhite\,(snow).$$

is of the general form

$$\text{(T)} \qquad \text{“}p\text{” is true iff } p.$$

called a T-sentence, which can be further specified as

$$\text{(T)} \qquad \text{“}p\text{” is true (in } \mathfrak{L}_1\text{) iff } p \text{ (in } \mathfrak{L}_0\text{).}$$

In particular, $\mathfrak{L}_0$ does not contain the sentence $IsTrue\,(p)$ (or $IsFalse\,(p)$, for that matter), nor the relations $IsTrue\,(x)$ or $IsFalse\,(x)$. This means that $\mathfrak{L}_0$ is a formal language that is *not closed semantically*, i.e. the theory of truth for $\mathfrak{L}_0$ is not construed within $\mathfrak{L}_0$.[8] It should now be obvious that any interpretation $\mathcal{I}$ for any language $\mathfrak{L}_0$ is always a construct of a metalanguage $\mathfrak{L}_1$. Hence, when one says that $Outcome\,(1,h)$ is true, one is actually expressing the T-sentence

$$\text{“The outcome of Flip 1 is heads” is true iff } Outcome\,(1,h).$$

[8]It is interesting to remark that a semantically closed language is host to the Liar Paradox: Consider the sentence "This sentence is false"; if it is true, then it is false, and if it is false, then it is true. It was precisely this and other semantic paradoxes that led A. Tarski to develop this purely formal analysis of truth (he called it the semantic theory of truth). The example given originally by Tarski (1944) has *snow is white* instead of $IsWhite\,(snow)$; he had formalized languages in general in mind, and not a specific (logical) language.

And, in effect, *Outcome* $(1,h)$ is a sentence of the language $\mathfrak{L} = \mathfrak{L}_0$, whose expressive capabilities are exhausted in expressions of this form (see Appendix).[9]

3.1.2 Facts and justifiers

Let it be granted that condition (i) of the tripartite analysis of knowledge has thus been satisfactorily settled, namely by means of the structure $\mathcal{M}_\phi = (\phi, 1)$ called a model for every fact ϕ. Then, the structure Ξ, in its simplest form, can be seen as a set of *true beliefs* held by one or more subjects S (say, the manager of Ξ and its users); hence, condition (ii) of the tripartite analysis is settled, too. It is now required that the same be done for condition (iii). Recall that Table 1 is the result of collecting data by using one's sensory modalities, vision in the case at hand: Every time the coin was flipped, an observer wrote down the number of the flipping event and the corresponding observed outcome. Is this epistemic justification enough? Unfortunately not, as our observer can be short-sighted, or they could be just inattentive for a myriad reasons (for instance, their mortgage payment was due shortly, or less plausibly but still possibly, there was a trickster demon interfering with their observations). Just as in the case of truth, what is required is a formal notion of epistemic justification; in particular, we need a formal theory of epistemic justification whose central assumption is that

Every fact has a justification.

I shall call this the *universal justification assumption* (UJA). In the following paragraphs, I sketch such a theory, leaving a more comprehensive elaboration for future work. (In the meantime, the reader is referred to the work of S. Artemov and colleagues–e.g., Artemov & Fitting, 2019; Artemov & Nogina, 2005–, some aspects of which this sketch–tentatively–draws on.)

To begin with, it makes sense to ask what formal structure can be conceived for epistemic justification of *beliefs in general.* We need a pair (ϕ, Jx), where ϕ is a belief and J is a binary relation between ϕ and some (reason) x such that x accounts for the truth (or validity) of ϕ or, in other words, is a *justification* or *proof* of ϕ, written $x : \phi$. Let us call this pair a *justifier for* S and denote it by $\mathcal{J}_S$. For any particular belief ϕ, S fails to know ϕ if $\mathcal{J}_S = (\phi, \emptyset)$: The pair $\mathcal{J}_S = (\phi, \emptyset)$ formalizes the case that ϕ is not a fact, or true belief. A justifier for S is more generally a pair $\mathcal{J}_S = (\mathcal{P}, Just)$ where $\mathcal{P} = \{p_1, p_2, ..., p_k\}$ is the set of propositions believed by S and $Just = \{Jx_1, Jx_2, ..., Jx_m\}$ is the set of corresponding justifications.[10] *Just* may be larger than $\mathcal{P}$ because there may be, and often there are, more than one justification for a single true belief, but the reverse can be the case, too, as different true beliefs may share justifications; in particular, some of–or, though uncommonly, even all–the beliefs in $\mathcal{P}$ may have no justification. If we make a justifer a more precise formal structure by adding the universe–or domain(s) thereof–over which S has beliefs, then

[9]This means that any *sentence* $\phi \in \mathcal{P}_{\mathfrak{L}}$ is safely assumed to be true, as we have it that $\phi = \neg\phi$ or $\phi = \neg(\neg\phi)$, by an application of double negation. A logical language whose set of sentences contains both ϕ and $\neg\phi$ gives rise to a trivial logical system, namely via inconsistency.

[10]More strictly from an algebraic viewpoint, and whenever formally required, we have the pair $\mathcal{J}_S = \left(\left(\mathcal{P} = \{p_i\}_{i=1}^{k}\right) \cup \left(\{x_j\}_{j=1}^{m} = Just\right), J\right)$, where $\mathcal{P} \cap Just = \emptyset$ and $p_i J x_j$ is a binary relation. This makes of $\mathcal{J}_S$ a *frame* in algebraic jargon.

these latter cases would be formalized as $\mathcal{J}_S^{\mathbf{U}} = (\mathcal{P}, Just \cup \emptyset)$ and $\mathcal{J}_S^{\mathbf{U}} = (\mathcal{P}, \emptyset)$, formalizing *partial knowledge* and *ignorance*, respectively.

It should be noted that S does not feature in the justifier for S: In effect, a justification relation ϕJx such that $x : \phi$ is completely independent from any belief holder; it is a feature of reality, to put it this way, one to which S find themselves in a relation by holding the belief that ϕ. Contrarily to the case in mainstream epistemology, the character of this relation is wholly irrelevant: S may believe that ϕ for the wrong reasons, or accidentally; what matters is that there is a justifier for S. For instance, and invoking a well-known illustration in mainstream epistemology (e.g., Kripke, 2011), seeing the façade of a red barn Henry (S) believes that it is indeed a barn that he is seeing. In fact, Henry has been driving along a country road where fake façades of blue barns have been strategically posited for tourists, only he is unaware of this. When seeing this particular barn–a real red barn–he truly believes it is a barn he is seeing, because he believes he has been seeing a lot of barns. According to some mainstream epistemologists, Henry is not justified in his belief, so he fails to know that it is a barn he is seeing. This paradoxical scenario is eliminated by introducing a structure $\mathcal{J}_{Henry} = (q, Jr)$ where q = "This is a barn (I'm seeing)," because we have qJr for r = "This is a real barn," and $r : q$, i.e. this reason justifies Henry's belief.

As said above, a true belief that ϕ may have several justifications; these may then be ordered by *strength* as $Jx_1 \preceq ... \preceq Jx_l$ where $Jx_i \preceq Jx_j$ is read "Jx_j is a stronger justification than Jx_i." For instance, suppose that some subject S believes they have acne (q) because they have spots; as a matter of fact, S has folliculitis, of which acne may be considered a special case. Given reasons s = "The spots are caused by folliculitis" and t = "The spots are caused by acne," we have $qJs \succeq qJt$, i.e. Js is a better (or stronger) justification than Jt for q, and thus $s : q \succeq t : q$. That S actually believes that $t : q$ weakens their epistemic justification, but it does *not* obliterate it: No matter how weak one's justification for ϕ might be, it contributes to knowledge of ϕ. The strongest justifier of all for any ϕ is the pair $\mathcal{J}_S = (\phi, \models)$ where $\models$, abbreviating $J \models$, denotes the binary logical consequence relation of the deductive kind.[11] This is followed by $\mathcal{J}_S = (\phi, |\approx)$, where $|\approx$ is some non-deductive logical consequence, i.e. mostly abductive, inductive, or probabilistic consequence. In particular, a justification $J = \models$ or $J = |\approx$ for ϕ can actually subsume, or entail, the truth of ϕ, thus greatly simplifying the analysis of knowledge (Artemov & Nogina, 2005; Augusto, 2011).[12] This can be obtained by postulating that for any sentence p and some proof t,

$$(Factivity) \qquad \text{if } t : p, \text{ then } p \text{ is true.}$$

In other words, justification is sufficient to conclude that a belief is true. As seen above summarily in the red barn and the folliculitis examples, the principle of *factivity* may also hold outside the realm of logic proper if, by following UJA, we succeed in formalizing ϕJr when r is a reason to be found in reality; this done, the metaphysical

[11] In this particular case, $\models: \phi$ is typically written as $\models \phi$. As a matter of fact, $\models$ can simply replace $J \models$, as it is a binary relation (we have $\emptyset \models \phi$ or $\mathcal{P}^{(')} \models \phi$) that is in fact a consequence justification in logic. See Augusto (2020d).

[12] In Augusto (2011), I do this subsumption from an informal, greatly mainstream approach that is at the same time a defense of pragmatism. This stance can be formalized in a general way for r = "It pays off to believe," so that ϕJr is read "It pays off to believe that ϕ."

battles over truth become irrelevant, as postulated by Tarski (1944). Actually, we not only simplify the analysis of knowledge, but also have now an axiomatizable system in which complex justifications can be formally computed. For instance, besides the tautologies of classical logic and the rule *modus ponens*, we may axiomatize factivity as

$$\text{(F)} \qquad t : p \rightarrow p.$$

Additionally, we may have the following axioms of justification for propositions ϕ, ψ, proofs s, t, and the binary operations *application* $(\cdot)$ and *sum* $(+)$, which may be added to $Op_{\mathfrak{L}}$ in a purely *ad-hoc* manner:

- (A) $s : (\phi \rightarrow \psi) \rightarrow (t : \phi \rightarrow [s \cdot t] : \psi)$
- (S) $s : \phi \rightarrow [s + t] : \phi; \quad t : \phi \rightarrow [s + t] : \phi$

Intuitively, axiom A (for *Application*) allows the construction of a proof $s \cdot t$ from both proofs s and t, and axiom S (for *Sum*) formalizes the monotonicity of computing with justifications: Once a proof, say s, has been found for ϕ, ϕJs remains a justification even if additional evidence is provided by $s + t$, i.e. we have $\phi J\,(s + t)$.

In any case, what we now have is a set $\mathbf{J} = \{\phi \mid t : \phi\}$ of *justified sentences* that is not dissimilar from the set $\mathbf{L} = \{\phi \mid\, \models \phi\}$ of *tautologies* that constitute a specific logic. Let $\mathbf{J}$ and $\mathbf{L}$ be considered as KBs; then, $\mathbf{L}$ is *more reliable* than $\mathbf{J}$ if $(\models : \phi) \succeq (t : \phi)$ for any proof t such that $t \neq \models$; otherwise, they are equally reliable KBs. We can, as a matter of fact, speak of *coincidence* of a justifier and a model for ϕ when ϕ is a tautology, a sentence that is true in all interpretations. Because facts correspond to sentences that are always true *at least* in some universe $\mathbf{U}$ (or some domain $\mathscr{D} \subseteq \mathbf{U}$) we can extend this coincidence to models and justifiers in general.[13]

Henceforth, I shall speak mostly of truth, but the reader should bear in mind UJA and this coincidence of models and justifiers.

3.1.3 Conditional facts and rules

The above can now be extended to *compound* (or *complex*) *propositions*. Let $p, q \in \mathcal{P}_{\mathfrak{L}}$ be atoms; then $\neg p$ or $\neg q$, $p \wedge q$, $p \vee q$, and $p \rightarrow q$ are well-formed formulas of $\mathfrak{L}$, namely compound formulas. The meaning of compound propositions is a function of the meaning of their constituents, a property that is formalized as follows:

- $val\,(\neg p) = 1$ iff $val\,(p) = 0$
- $val\,(p \wedge q) = 1$ iff $val\,(p) = 1$ and $val\,(q) = 1$
- $val\,(p \vee q) = 1$ iff $val\,(p) = 1$ or $val\,(q) = 1$
- $val\,(p \rightarrow q) = 1$ iff either $val\,(p) = 0$ or $val\,(q) = 1$

[13] Note that whenever ϕ is a tautology we write simply $\models \phi$ (compare with $\models_{\mathcal{M}} \phi$ above). From an algebraic perspective, both $\mathcal{J}_S = (\phi, J)$ and $\mathcal{M}_\phi = (\phi, 1)$ are *frames*, models whose relation R is a binary relation. As seen, both J and 1 can be realized by the binary relation $\models$. Then, we can settle the identity $\mathcal{J}_S = (\phi, \models) = \mathcal{M}_\phi$ in general for tautologies and the identity $\mathcal{J}_S^{\mathbf{U}} = (\phi, \models) = \mathcal{M}_\phi^{\mathbf{U}}$ specifically for some domain $\mathbf{U}$.

This can be extended by induction to any formulas $\phi, \psi \in \mathcal{P}_{\mathfrak{L}}$.

In terms of epistemic justification, and so as to make this coincide with truth by invoking the factivity principle, we have then for some proof t–and greatly, if not grossly, simplifying justification for compound propositions, especially in the cases of $\neg p$ and $p \to q$ (but see below):

- $t : \neg p$ iff $\emptyset : p$
- $t : (p \wedge q)$ iff $t : p$ and $t : q$
- $t : (p \vee q)$ iff $t : p$ or $t : q$
- $t : (p \to q)$ iff either $\emptyset : p$ or $t : q$

We speak now of *complex facts* and a KB $\Xi = \{p_i\}_{i=1}^{k}$ where at least one of the p_i is a complex fact may be called a *complex KB* (more often than not just *KB*).

Conditional facts, i.e. complex facts of the form $p \to q$ and so called because the operator $\to$ expresses a condition ("if p, then q") between an antecedent p and a consequent q, require special attention: In order for a true proposition like $p \to q$ to be a conditional fact it must be the case that $p \to q$ be read as "if p is true, then q is true," or equivalently "either $\neg p$ or q is true." Otherwise, sentences such as "If the moon is made of cheese, then all humans are mortal (immortal)," where the antecedent is false and the consequent is true (false, respectively), would be facts. This remark is important, because in classical logic, as seen above, from a false antecedent anything follows, be it truth or falsity. This classical principle, technically called *ex falso quodlibet*, cannot hold insofar as facts are concerned. In the metalanguage of $\mathfrak{L}$, a conditional proposition $p \to q$ is a fact iff $\models p \to q$, i.e. $p \to q$ is a tautology.[14]

A conditional fact often has the form

$$B_1 \wedge ... \wedge B_k \to A$$

where $B_1, ..., B_k, A$ are all true atoms of the form $R(c_1, ..., c_n)$, the $\{c_j\}_{j=1}^{n}$ are all constants from a specific universe. For instance, the proposition

(q) If John is older than Mary and Tessa is older than John, then Tessa is older than Mary.

formalizable in $\mathfrak{L}$ as

$$(q) \qquad Older\,(john, mary) \wedge Older\,(tessa, john) \to Older\,(tessa, mary)$$

[14]We have $\models p \to q$ essentially when $p = q$. This, however, if far from trivial, as it is often the case that the identity, or equivalence, between p and q is not obvious. For instance, "If the patient has pyorrhea, then he has periodontitis" and "If they saw the morning star or the evening star, then they saw Venus" are both tautologies, because in fact pyorrhea and periodontitis designate one and the same clinical condition, and the morning/evening star just is another name for the planet Venus. This, in particular the latter example, poses interesting philosophical questions concerning sense and reference with relation to knowledge that are beyond the scope of this article, but see Frege (1892) for a well-known philosophical discussion on sense and reference.

where *Older* abbreviates *IsOlderThan*, is a conditional fact. Importantly, the *deduction theorem* (DT) assures us that if all the B_i are true and A is true, then $B_1 \wedge ... \wedge B_k \rightarrow A$ is a tautology, i.e.[15]

$$\text{(DT)} \qquad \text{If } (B_1 \wedge ... \wedge B_k) \models A, \text{ then } \models (B_1 \wedge ... \wedge B_k) \rightarrow A.$$

Conditional facts are important knowledge structures, because they are associated with other knowledge structures called *rules*. What distinguishes a rule from a conditional fact is the presence of individual variables in the former. In particular, in a rule $B_1, ..., B_k, A$ are all atoms of the form $R(t_1, ..., t_n)$ where at least one of the terms $t_j \in ((\bigcup B_i) \cap A)$ is an individual variable, the B_i are called *goals*, and A is a *fact*. The rule form is actually more often than not formulated as

$$(r) \qquad A \leftarrow B_1, .., B_k$$

where A is called the *head* (of the rule) and the B_i constitute the *body* (of the rule). Thus, a rule corresponding to proposition q can be formulated as

$$(r_q) \qquad Older\,(x, y) \leftarrow Older\,(z, y)\,, Older\,(x, z)\,.$$

A (complex) KB is accordingly more often than not a pair

$$\Xi = \left(\{p_i\}_{i=1}^{k}, \{r_j\}_{j=1}^{m} \right)$$

where the p_i are (typically) atomic facts and the r_j are rules.[16] Although the body of a rule r is said to be composed of goals, these are also facts. Indeed, if there are no concrete individuals to replace the variables in $Older\,(z, y)$ and $Older\,(x, z)$, so that these two atoms are facts, rule r_q is useless. This is formally secured by what can be called the epistemic equivalent to the deduction theorem, the *knowledge generation theorem* (KGT):[17]

$$\text{(KGT)} \qquad \text{If } B_1^{'}, ..., B_k^{'} \text{ are facts and } A' \leftarrow B_1^{'}, ..., B_k^{'} \text{ is a fact, then } A' \text{ is a fact.}$$

Informally expressed, we have it that *knowledge is solely generated from knowledge.* We can see this as an intuitive formulation of *correctness* for a KB Ξ. If additionally all the rules of a KB Ξ together with its facts are sufficient to *generate all the knowledge* contained in Ξ, then Ξ is said to be *complete.* A KB Ξ is said to be *adequate* if it is both correct and complete.

[15] Alternatively, $\{B_i\}_{i=1}^{k} \models A$ iff $\{B_i\}_{i=1}^{k-1} \models B_k \rightarrow A$, $\{B_i\}_{i=1}^{k-2} \models B_{k-1} \rightarrow (B_k \rightarrow A)$, and so forth until we have $\emptyset \models B_1 \rightarrow (... \rightarrow (B_{k-1} \rightarrow (B_k \rightarrow A)))$.

[16] There are different configurations for this pair. For instance, both sets of Ξ are stored together in Prolog, but separately in Datalog, in which $\{p_i\}_{i=1}^{k}$ is a KB proper and $\{r_j\}_{j=1}^{m}$ is a *program.* See Augusto (2020a, Chapter 9).

[17] In logical terminology, the atom $A' = R(c_1, ..., c_n) = R(\overrightarrow{c})$ where $\{c_i\}_{i=1}^{n}$ are constants, is called an *instance* of $A = R(x_1, ..., x_n) = R(\overrightarrow{x})$. Then, A' is a fact iff there is some interpretation $\mathcal{I}_{\overrightarrow{c}}^{\overrightarrow{x}} = \mathcal{I}$ such that $val_{\mathcal{I}}(A') = 1$. This, in turn, is only the case if there are instances $\left\{B_i^{'}\right\}_{i=1}^{k}$ of the $\{B_i\}_{i=1}^{k}$ such that $A' \leftarrow B_1^{'}, ..., B_k^{'}$ is an instance of some rule $A \leftarrow B_1, ..., B_k$.

3.2 Dissolving the DIK Trinity. II: Knowledge vs. Data; Knowledge vs. Information

This–facts, rules, and their assembling into knowledge bases–being established, it is now required that knowledge be distinguished from both data and information, so that the DIK trinity be fully dissolved.

Knowledge vs. Data – Let us consider an object of the form p: p is either a fact or a datum, according to whether it is interpreted by means of some interpretation $\mathcal{I}$ such that there is a model for it, or no interpretation $\mathcal{I}$ is considered, respectively. As seen above, in constructing a KB Ξ a segregation is required between an object language, say $\mathfrak{L}_0$, constituted by sentences, and a metalanguage, say $\mathfrak{L}_1$, such that a sentence $p \in \Xi$ iff a corresponding T-sentence can be stated. No such considerations are required for a database Δ. Furthermore, given a database Δ, there is no need for a justifier of the form $\mathcal{J}_S = (\phi, Jx)$, as Δ does not correspond to a set of justified true beliefs held by some subject S, simply because no effort whatsoever need be made in that sense. As a matter of fact, it is difficult for humans to construct a database, as epistemic justification and truth are more often than not almost intrinsic concerns for human agents, not the least reason for this being the "intuition" that without knowledge our well-being is far from secured.[18] But a satellite collecting data does precisely that: It constructs a database Δ. This database Δ may then become a KB Ξ if there is some human agent (or some other belief-holding agent) with respect to whom/which all the data in Δ are, or can become, justified true beliefs.

Figure 4 shows this distinction between knowledge and data.

Knowledge vs. Information – Above, it was seen that $Outcome\,(1,h)$ is a sentence of our KB Ξ_{Flip}. Then, presumably it is the case that there is a justification of the form $t : Outcome\,(1,h)$, and we assume that there is a justification structure $\mathcal{J}_S = (Outcome\,(1,h)\,, Jt)$. This, by UJA and the principle of factivity, assures us that we know that the outcome of Flip 1 was heads. However, if the outcome of Flip 1 was in fact tails, we have the structure $\mathcal{J}_S = (Outcome\,(1,h)\,, \emptyset)$, and we have no knowledge with respect to this particular flip. After all, and as emphasized above, the observer might have been wrong when noting down the outcome of Flip 1. This is precisely where *knowledge* distinguishes itself from *information*: When the epistemic status of our beliefs (or of the facts in a KB) is uncertain or indeterminate, our facts become bits. Just as in the case of encoding data, entropy characterizes our beliefs in the sense that for every proposition p that we believe, it is the case that either p is true or p is false.[19] Then, just as for the possible outcomes when flipping a coin, if we consider p as a random variable we have $\Pr\,(p = x_i) = 0.5$ for $x_1 = true$ and

[18] In fact, our propositional-like conceptions of the world and their corresponding concepts can be seen as facts, so that our terminological constructs may in fact correspond to a knowledge base. See Badie (2020a; this issue) for an account of this hypothesis from the viewpoint of constructive epistemology.

[19] This holds even for valuations in many-valued logics, as every many-valued logic can be reduced to bivalent logic by means of a structure $\mathfrak{M}$ called a *matrix* that allows us to "send" all formulas valuated within a set of distinguished values $D \in \mathfrak{M}$ to a set of true formulas, all the remaining formulas being considered *not-true* (rather than false). This "sending" is actually a homomorphism, and matrix theory is elaborated on in algebraic terms; as this is outside the scope of this paper, I refer the reader to Augusto (2020c) for a comprehensive discussion, or to Augusto (2020d) for a briefer discussion.

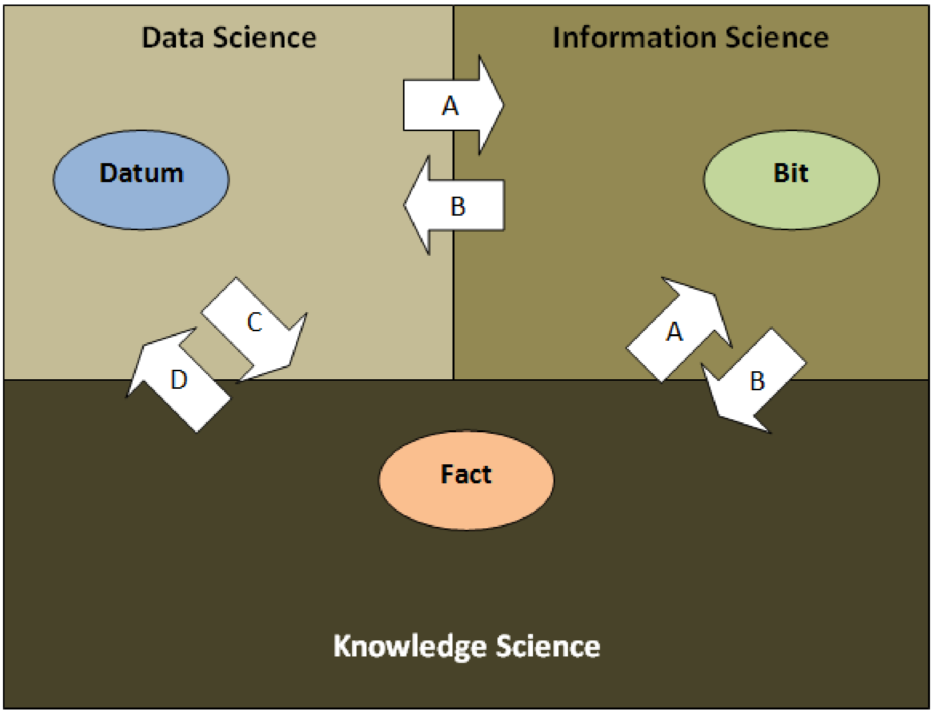

Figure 4: Data vs. Knowledge vs. Information. A: $H(X) > 0$; B: $H(X) = 0$; C: $val_{\mathcal{I}}(X) \in \{1\} = (X, 1) = \mathcal{M}_X$; D: $val_{\mathcal{I}}(X) \in \emptyset$.

$x_2 = false$. Thus, the information contained in any belief that p is exactly equal to 1 bit, as we have

$$H(p) = -\left(\left(\frac{1}{2}\log_2\frac{1}{2}\right) + \left(\frac{1}{2}\log_2\frac{1}{2}\right)\right) = 1$$

and all the computations in information theory can be applied to beliefs; for instance, we can compute the mutual information of two beliefs p and q. A KB becomes then an *information repository* (just to use an usual label). In the absence of entropy with respect to beliefs, we are back to a KB. See Figure 4.

The key word here that allows us to distinguish a KB Ξ from a database Δ is, as already mentioned above, *belief*: If X is taken as a mere random variable such that $H(X) = 0$, then it is a datum, and any collection of data is a database; if, however, X is taken as a belief, then $H(X) = 0$ indicates that X is a fact and any collection of facts is a KB. Whenever $H(X) > 0$, we are in the realm of information, in which it does not really matter whether one is dealing with facts or mere data: There is *uncertainty* with respect to either.

An illustration: On personal identification cards, it is usually the case that the age of the card holder is shown. Let us consider this to be expressible by the string $IsOld(x, y)$ where x stands for the individual and y for the age in years. Then, say,

$IsOld\,(john, 64)$ is a datum, i.e. a string of symbols resulting from an empirical collection (John was asked how old he was by some public servant, or his age was verified in some archive, etc.). If $IsOld\,(john, 64)$ is a justified true belief (e.g., the archive is highly reliable and John is indeed 64 years old), then there is an interpretation $\mathcal{I}$ that makes this datum into a fact. Regardless of whether $IsOld\,(john, 64)$ is a datum or a fact, it becomes information when, for instance, an employer considers that John might be too old for the job, or his physician considers whether to increase some nutritional supplement intake.

In other words, and to sum up, both data and facts are inert as far as action is concerned; every action entails a degree of uncertainty, and that is when data or facts become pieces of information. The difference between data and facts is extrinsic to action and is a purely formal one. The point then is to be certain that some piece of information is indeed a fact, so as to secure *rational action* (when the agent is a human), or the *right action given the environment* (for instance, in the case of a robot). Such certainty is purely ideal, as it will be seen, but it sets the standards for the design and construction of knowledge systems.

4 Knowledge Systems

KBs are typically not static structures: Besides exhibiting internal behavior of some sort (e.g., by means of a fully automated calculus; see Augusto, 2020a), they typically require more or less frequent external actions, namely updating and maintenance. Although largely task-independent, they are designed with specific aims in view, often the solution to some particular (class of) problems.[20] To these ends, they require both processes to be carried out over them and agents that carry them–all or a part of them–out, i.e. they are components of a knowledge system.

A *knowledge system* (KS) is a triple

$$\mathbb{K} = (\mathscr{KS}, \mathscr{KP}, \mathscr{KA})$$

where $\mathscr{KS}$ is a collection of *knowledge structures*, $\mathscr{KP}$ is a collection of *knowledge processes*, and $\mathscr{KA}$ is a collection of *knowledge agents*. Generally considered, namely from a more behavioral perspective of both cognitive science and computer science, a KS is some structure that upon knowledge-structured input outputs an appropriate action. As already stated above, when the knowledge agent is a human, this is expected to be a rational action; for artificial agents one speaks of the right action given the environment. In either case, it is assumed that only facts, atomic or complex, can lead to the rational or the right action. Hence, facts, taken as justified true beliefs as

[20]This holds particularly at the ontological level. For instance, facts of the form $Teaches\,(x, y, z)$ where x stands for some faculty member, y does so for a course, and z for the respective level, are more appropriate for a KB of instructors than $Takes\,(x, y, z)$ where x stands for a student and y, z are as above. In the first case, one may be interested in finding out who is teaching what at which level, whereas in the second the problem might be to find out who is taking what at which level. But minor as this distinction of domains might be, it already requires a good grasping of ontological constructs. In effect, ontology is an essential component of knowledge systems, a topic which is beyond the scope of this text, but see Limbaugh et al. (2020; this issue) for an account of how ontology is associated with cognition in "intelligence systems" seen as knowledge systems, and see Saba (2020; this issue) for some relations between ontology and natural language from the viewpoint of commonsense and/or background knowledge.

settled above, are the structures that account for the invariance of the term *knowledge* in the three collections above: $\mathscr{KS}$ is a collection of facts, $\mathscr{KP}$ is a collection of processes over facts, and $\mathscr{KA}$ is a collection of agents in possession of, or looking for, facts. Thus, $\mathbb{K}$ can be realized in a plethora of systems–expert systems, neural-network systems, intelligent systems, organizations, ...–including the human brain, as long as they, being *input facts* by an agent (which has the responsibility of turning data into facts), *output rational, or the right, actions.* Figure 5 shows the general schema of a KS $\mathbb{K}$; below, the components of this system are elaborated on.

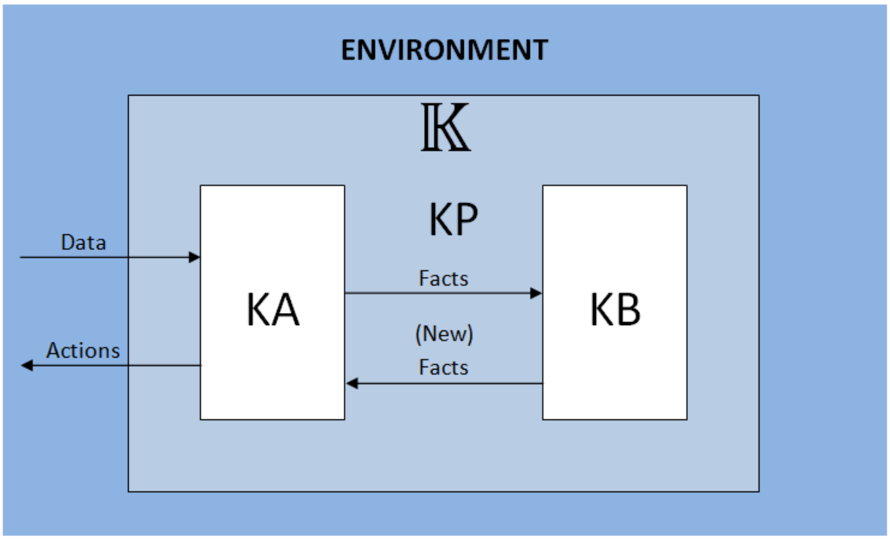

Figure 5: A general KS.

Strictly defined, a KS $\mathbb{K}$ is a structure that generates new knowledge, i.e. *new facts.* This is then the restricted case when action is the output of facts. I shall focus on *rule-based* KSs, triples of the form $\mathbb{K} = \left(\overrightarrow{\Xi}, O, A\right)$, where $\overrightarrow{\Xi} = \Xi \cup Rul$ is a *rule-based* KB, $Rul = \{r_i\}_{i=1}^{n}$ is a set of rules of the form $A \leftarrow B_1, ..., B_k$, $O = \{o_j\}_{j=1}^{m}$ is a set of operations or processes on $\overrightarrow{\Xi}$, and $A = \{a_i\}_{i=1}^{k}$ is a set of agents.[21] For simplicity, I consider only $|A| = 1$. In particular, I shall be concerned with systems such that $O' = \{\models\} \subseteq O$, often called *deduction systems.* They are also called *goal-directed* systems because they use *backward-chaining reasoning*, i.e. they reason from goals to facts. In these systems, a rule has a strictly *logical* or *declarative* reading (cf. DT and KGT): If all the B_i are true (or can be proved), then A is true (or can be proved). Diagnostic systems and automated theorem provers are examples of deduction systems.[22]

[21]The prototypical example of a KB $\overrightarrow{\Xi}$ is a Prolog program, in which $\Xi \cup Rul$ is stored in the same *locus.*

[22]The other are called *reactive systems* and they often use *forward-chaining reasoning*: From input facts they reason toward conclusions. In these, a rule has a *procedural* interpretation: If B_1 and ... and B_k are the case, then DO A. Intelligent systems that interact with the environment are typically reactive systems.

4.1 Knowledge Processes

Although what I discuss next can hold for any KS generally taken (see Fig. 5), it holds in particular for KSs as artificial intelligence (AI) constructs known as computer programs. KSs require a few mundane actions be carried out, namely design, construction, manipulation, and maintenance, but these are not the processes I shall consider as *knowledge processes* (KPs) proper. Given a KB Ξ, there are four main KPs over Ξ: *Representation*, *Acquisition & Storage*, *Generation*, and *Access & Share*.

In a KS of the deductive kind $\mathbb{K} = \left(\overrightarrow{\Xi}, O, A\right)$, these four processes actually correspond to the set $O = \{\texttt{tell}, \texttt{ask}\} \cup \{\models\}$, where $O'' = \{\texttt{tell}, \texttt{ask}\}$ contains the basic operations that are in general common to KSs: `tell` is for adding new facts to the system (and obviously storing them), which presupposes a representation medium (a language), and `ask` for finding out (and sharing) what the system knows. For the sake of generality, I assume that the *representation language* of the KB, which is here $\mathfrak{L}$, and the *interaction language* of the agent are essentially the same, merely with the extension $Op_{\mathfrak{L}} \cup \{., ?\}$ such that "." at the end of a well-formed string $\phi \in \mathcal{P}_{\mathfrak{L}}$ indicates that ϕ is a fact, and $\phi?$ indicates a query to the KB. (For simplicity, I write "ϕ." as ϕ simply.) The operation $\{\models\} = O'$–which is in fact a *relation* but is here taken in its natural association with the consequence *operation* (see Augusto, 2020d)–is specific to deduction KSs. In this Subsection, I concentrate on the processes in which the agent can be omitted (Representation) and which are internal to the KB (Generation). In the next Subsection, I address the processes that require an agent, to wit, the two remaining processes of the quadruple above.

Knowledge Representation – Knowledge can, in principle, be represented in many media (e.g., graphs, frames, ...). If one's choice falls on a formal language, then this should be as high-level and as expressive as possible, a feature that has to be weighed against its computational costs, in particular for logical languages. Let us consider this to be the *symbolic level* of a KS. With respect to the first aspect, $\mathfrak{L}$ can be used to express virtually not only all assertions, but also all corresponding queries in any natural language. This is accounted for by the ability of this language to express such minor changes in meaning as in the propositions "It *tastes like* vanilla." and "It *likes* vanilla *taste*.", which can be existentially formulated as $\exists x \exists y\, (TastesLike\,(x, y))$ and $\exists x \exists y\, (LikesTaste\,(x, y))$. Given these existential formulations, one can make queries such as $TastesLike\,(icecream, vanilla)?$ and $LikesTaste\,(fritz, vanilla)?$ for the very large domains corresponding to the different subjects expressed by the natural language pronoun "it." With respect to its computational features, $\mathfrak{L}$ is essentially decidable, as it does not have either function symbols or the identity symbol (=). This said, it is not cost-free in terms of computational complexity; as a matter of fact, it is still so costly that the trade-off between expressivity and complexity may be negative on the latter side.[23]

But when choosing a representation language for knowledge one's consideration should not only fall on the symbolic level of the system, but also–or especially–on what Newell (1982) called the *knowledge level*: A representation medium must not

[23] The decidability of $\mathfrak{L}$ depends on other aspects, in particular on electing Herbrand semantics, by means of which the "propositionalization" of a first-order language can be obtained. See Augusto (2020a, c). For the trade-off between expressivity and complexity for (an extension of) the standard first-order language, see Levesque & Brachman (1987), an early discussion of this topic.

only be able to *represent something* (an object, a relation, a process, a state, ...), but it must also *embody the knowledge of the system about that thing.* As seen above, a construct of the kind $R(t_1, ..., t_k) \in \mathcal{P}_{\mathfrak{L}}$, known as an atom, does not only represent something (e.g., $IsOlderThan(x, y)$ represents the age relation between some unspecified individuals x and y), but if $val_{\mathcal{I}}(R(c_1, ..., c_k)) = 1$ for some interpretation $\mathcal{I}^{\vec{x}}_{\vec{c}} = \mathcal{I}$, then this atom actually embodies the knowledge about this thing in an atomic fact (e.g., $IsOlderThan(john, mary)$.). The same holds for constructs of the form $A \leftarrow B_1, ..., B_n$, known as rules, where each of $A, B_1, ..., B_n$ is of the form $R(x_1, ..., x_k)$: If $val_{\mathcal{I}}(\bigwedge_{i=1}^{n} B_i) = 1$ and $val_{\mathcal{I}}((\bigwedge_{i=1}^{n} B_i) \rightarrow A) = 1$, then $val_{\mathcal{I}}(A) = 1$ for some interpretation $\mathcal{I}^{\vec{x}}_{\vec{c}} = \mathcal{I}$, and we have both a complex fact and a new fact. (Recall: *Ex falso quodlibet* is excluded with relation to facts.) This latter case entails that the KB embodies more knowledge than that which is explicit in it. In this sense, the purely formal concept of *interpretation* does indeed become an epistemic concept: An interpretation of atoms and rules yields a representation of facts that go on to embody the knowledge of a KS in a KB.

Still with relation to the knowledge level, if one extends $\mathfrak{L}$ with the unary operator K, read "it is known," then one can reason about the knowledge embodied in a KS $\mathbb{K}$ and the relation of agents to their own knowledge in the system. For instance, the rule $(p \rightarrow Kp)^{\mathbb{K}}$ expresses the case that if p is a fact in $\mathbb{K}$ then p is known in $\mathbb{K}$, and $(K_a p \rightarrow K_a K_a p)^{\mathbb{K}}$ expresses the axiom that if an agent a knows p in $\mathbb{K}$ then a knows (in $\mathbb{K}$) that they know p in $\mathbb{K}$. This extension, known as *epistemic logic*, not only keeps all the classical tautologies expressible in $\mathfrak{L}$, but also has more tautologies; this shows its increased expressive power (unfortunately at the cost of computational efficiency; see above).

Finally, only a logical language gives us the assurance that our facts correspond not only to true propositions, but also to justified beliefs; this it does by means of the metalanguage operator $\models$. (One can also argue that human thinking is essentially logical–e.g., Augusto, 2014–but I am leaving this out here.)

Knowledge Generation – Although very simple, Ξ_{Flip} allows the extraction of knowledge by means of queries (the operation `ask`) such as $Outcome(1, t)?$, which will output the result "*no*", and $Outcome(x, t)?$, outputting $x = 2, 3, 4, 7, 8$. However, none of the replies we can obtain from this KB will *generate new knowledge* or, more correctly, *new facts*. In order to obtain this we need rules. An illustration: Let us suppose that we have a KB of rescued birds in an avian center with facts such as $IsBird(hen, lolita)$., $IsBird(turkey, peter)$., $IsBird(canary, roberto)$., etc. (Fig. 6 shows the complete KB in a Prolog implementation.) Although rescued, not all the birds in the center are sick, in which case they are quarantined and accordingly entered in the KB $\vec{\Xi}_{Birds}$ as the fact $IsQuarantined(x)$. (for instance, $IsQuarantined(bob)$.). Then, in order to find out which birds are sick it suffices to design the rule

$$(r_1) \qquad IsSick(y) \leftarrow IsQuarantined(y) .$$

Note that the fact $IsSick(bob)$. is not *explicitly* in the KB; the difference here between the operation `tell`, by means of which new facts can be added to the KB, is that $IsSick(bob)$. is already *implicitly* in the KB, and thus it does not require the operation `tell` to be carried out over the KB, namely by an agent outside the KB. In effect,

$IsSick\,(bob)\,.$ is an *inferred fact*, whereby it is meant

$$\overrightarrow{\Xi}_{Birds} \models IsSick\,(bob)\,.$$

and a set of operators $O \supseteq \models$ is accordingly called an *inference engine* (over a KB $\overrightarrow{\Xi}$).

New facts are essential to output further new facts. For instance, suppose that staff in the avian center are required to know which birds are sick, in order to provide them with a special diet of their natural food. These new facts can be generated by means of the rule

$$(r_2) \qquad IsOnDiet\,(y,z) \leftarrow IsBird\,(x,y)\,, sick\,(y)\,, eats\,(x,z)\,.$$

By means of this rule, the new facts $IsOnDiet\,(roberto, seeds)\,.$ and $IsOnDiet\,(bob, all)\,.$ can be generated, as in fact we have, for instance,

$$\overrightarrow{\Xi}_{Birds} \models IsOnDiet\,(bob, all)\,.$$

```
bird (penguin, toto) .
bird (ostrich, sheila) .
bird (emu, tom) .
bird (turkey, sam) .
bird (turkey, sandra) .
bird (hen, lolita) .
bird (canary, roberto) .
bird (nightingale, sarita) .
bird (crow, bob) .
bird (woodpecker, lola) .
bird (duck, cassandra) .
bird (duck, samantha) .
quarantined (roberto) .
quarantined (bob) .
eats (penguin, fish) .
eats (ostrich, all) .
eats (emu, all) .
eats (turkey, seeds) .
eats (hen, all) .
eats (canary, seeds) .
eats (nightingale, seeds) .
eats (crow, all) .
eats (woodpecker, bugs) .
eats (duck, all) .
sick (Y) : −quarantined (Y) .
on_diet (Y, Z) : −bird (X, Y) , sick (Y) , eats (X, Z) .
```

Figure 6: The KB $\overrightarrow{\Xi}_{Birds}$. (Source: Augusto, 2020a.)

We are assured that this, read as "the fact $IsOnDiet\,(bob, all)\,.$ holds in, or is entailed by, $\overrightarrow{\Xi}_{Birds}$", is indeed epistemically the case (i.e. it is a justified true belief), because a *proof* $t = \vdash$ can be produced such that we have

$$\overrightarrow{\Xi}_{Birds} \vdash IsOnDiet\,(bob, all)\,.$$

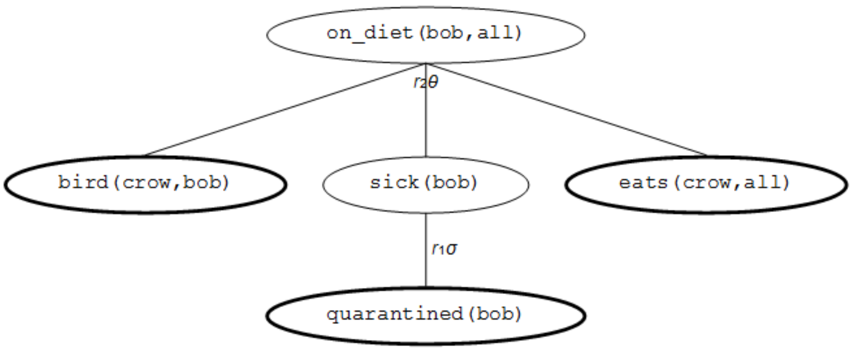

Figure 7: Proof tree for the inferred fact $IsOnDiet\,(bob)$. (Source: Augusto, 2020a.)

where the metalanguage symbol $\vdash$ denotes that the atom $IsOnDiet\,(bob, all)$. is *derivable* from $\overrightarrow{\Xi}_{Birds}$ by means of a logical calculus–say, the resolution calculus–such that a proof tree can be produced (see Fig. 7; σ denotes the substitution of *bob* for *y*, and θ the simultaneous substitution *crow* for *x* and *all* for *z*). In effect, the justifier $\mathcal{J}_S = (\phi, \models)$ for a sentence ϕ is alone no longer a sufficient justification, as the metalanguage relation $\models$, now also seen as a consequence operation, can output fewer facts than those entailable by the KB $\overrightarrow{\Xi}_{Birds}$; in particular, it might be the case that neither ϕ nor $\neg\phi$ is output. In other words, the KB may be *incomplete*. In order to make a KB *complete*, we need to pair it with a logical calculus, implemented by the inference engine, such that all the facts in it, and only the facts in it, are provable. This justifier can be denoted by $\mathcal{J}_S = (\phi, \models_\vdash)$ where $\vdash$, in turn, does not prove falsities (i.e. $\vdash$ is a *sound* relation of logical consequence). This, formalizable as

$$\overrightarrow{\Xi} \models \phi \quad \text{iff} \quad \overrightarrow{\Xi} \vdash \phi$$

provides us with an *adequate* KB $\overrightarrow{\Xi}$. Given this equivalence, we write henceforth simply $\models$ to denote an adequate consequence relation/operation.[24]

4.2 Knowledge Agents

Generally, a *knowledge agent* (KA) acting in an environment E is a triple

$$\mathcal{KA}_E = (\Xi_E, Goals_E, Actions_E)$$

where Ξ_E is the KB containing the facts relative to E, $Goals_E = \{g_i\}_{i=1}^m$ is a set of goals with respect to the environment E to be attained, and $Actions_E = \{act_j\}_{j=1}^l$ is a set of actions over E to be selected by the KA to pursue their goals in E. As already seen above, the agent acts in the environment according to their goals by relying on the facts in the KB; this can be generally called a *rational agent*. I leave here E largely undefined as "the world." In this, a KA can have goals as diverse as finding line

[24] More correctly, it is the logical system $\mathrm{L} = (\mathfrak{L}, \models)$ that is said to be adequate with respect to $\overrightarrow{\Xi} \subseteq \mathcal{P}_{\mathfrak{L}}$, which is being here considered as a theory. See Augusto (2020d).

connections in a metro network or diagnosing causes behind rare symptoms. Indeed, for every goal–or problem–there is, or can be in principle, a KB by means of which the KA can act rationally.

I omit here any considerations with respect to E, namely as far as data collection and actions are concerned, and consider a KA as a black box restricted to inputting facts into, and querying, a KB $\overrightarrow{\Xi}$; this agent can be given by the triple $\mathcal{KA}_{\overrightarrow{\Xi}} = \left(\overrightarrow{\Xi}, Goals_{\overrightarrow{\Xi}}, Actions_{\overrightarrow{\Xi}}\right)$. $KA_{\overrightarrow{\Xi}}$ is so restricted that we have $Actions_{\overrightarrow{\Xi}} = \{\texttt{tell}, \texttt{ask}\}$, where in fact $\{\texttt{tell}, \texttt{ask}\} = O_{\overrightarrow{\Xi}} \cap Actions_{\overrightarrow{\Xi}}$. This KA is solely expected to both "tell" the KB (and obviously store in it) the facts they have with respect to E and "ask" the KB about the–possibly new–facts it possesses concerning E (often with the aim of sharing them). I discuss these two operations from the viewpoint of generality.

Henceforth, I abbreviate an (adequate) KB $\overrightarrow{\Xi}$ as Ξ.

Knowledge Acquisition & Storage – Formally, the operation `tell` is a mapping

$$\texttt{tell} : \mathcal{P}_{\mathfrak{L}} \longrightarrow \mathcal{P}'_{\mathfrak{L}} \cup \mathcal{P}''_{\mathfrak{L}}$$

for $\mathcal{P}'_{\mathfrak{L}}$ the subset of sentences of $\mathcal{P}_{\mathfrak{L}}$ that are facts and $\mathcal{P}''_{\mathfrak{L}}$ the subset of formulas of $\mathcal{P}_{\mathfrak{L}}$ that are rules, by means of which a KA *tells* the KB Ξ that sentence ϕ is a fact or a rule. To simplify, I shall consider that $\mathcal{P}''_{\mathfrak{L}} \subseteq \mathcal{P}'_{\mathfrak{L}}$, by seeing rules as uninstantiated facts. The practical aspect with respect to the operation `tell` is that it increases a KB by adding new facts to it. An additional operation, not shared with the processes of Access & Share and to be applied only once, is the following:

$$\texttt{begin}\,() = \Xi^0.$$

It is important to remark that $\Xi^0 \neq \emptyset$: Ξ^0 contains all the *rules of inference* (and possibly *axioms*) of a calculus and perhaps tautologies of a specific logical system, i.e. consequences of the form $\models \psi$. In effect, Ξ^0 is the inference engine, which contains no facts about the environment but is solely designed with the operation $\models$ in view. The addition of a fact about the environment ϕ to Ξ^0 by means of `tell` is given by

$$\texttt{tell}\left(\phi, \Xi^0\right) = \Xi^1 = \Xi^0 \cup \{\phi\}$$

so that for any $i > 1$ we have

$$\texttt{tell}\left(\phi, \Xi^{i-1}\right) = \Xi^i = \bigcup_{i}^{i-1} \Xi \cup \{\phi\}$$

and obviously $|\Xi^i| = |\Xi^{i-1}| + 1$, where $|\Xi^i|$ denotes the number of *explicit* facts (which include uninstantiated rules) in Ξ^i.[25]

[25] Levesque & Lakemeyer (2000), which provides a comprehensive elaboration on the operations `tell` and `ask` (as well as on additional operations; see Chapter 5), specifies this operation as $\texttt{tell}\,(\phi_i, e_{i-1})$, where ϕ_i are the sentences to be added to the KB and e_{i-1} is the *epistemic state* corresponding to the KB before this addition. This agrees with their adoption of Kripke semantics, as is usual in the epistemic logic literature. I simplify by making each Ξ^i for $i > 0$ correspond to what the KB "knows" after the addition of a new fact. For these "epistemic states" taken in a rather loose sense, old Tarskian semantics should be enough. See Augusto (2019; 2020a, d).

Despite the formal definition of the operation `tell`, the cardinality of Ξ^i is, in principle, not a reliable measure of the knowledge contained in Ξ^i. In other words, one cannot say that a KB with, say, 100 propositions does correspondingly have 100 facts. To begin with, Ξ^i might contain *n implicit* facts not factored in in $|\Xi^i|$; additionally, as a result of the operation `tell`, it is often the case that *both* $\phi \in \Xi^i$ *and* $\neg\phi \in \Xi^i$ (the KB is inconsistent), or *neither* $\phi \in \Xi^i$ *nor* $\neg\phi \in \Xi^i$ (the KB is incomplete). These two latter scenarios are the result of misapplications of the `tell` operation: For each atom ϕ, either $\texttt{tell}\left(\phi, \Xi^{i-1}\right)$ or $\texttt{tell}\left(\neg\phi, \Xi^{i-1}\right)$, and at least one of $\texttt{tell}\left(\phi, \Xi^{i-1}\right)$ or $\texttt{tell}\left(\neg\phi, \Xi^{i-1}\right)$, should always be the case.[26] To sum up, the operation `tell` can cause the KB to be in one of four epistemic states with respect to a proposition ϕ (see Table 2).

`tell`	**Epistemic state**
$\texttt{tell}\left(\phi, \Xi^{i-1}\right)$	Ξ^i knows ϕ
$\texttt{tell}\left(\neg\phi, \Xi^{i-1}\right)$	Ξ^i knows $\neg\phi$
both $\texttt{tell}\left(\phi, \Xi^{i-1}\right)$ and $\texttt{tell}\left(\neg\phi, \Xi^{i-1}\right)$	Ξ^i knows both ϕ and $\neg\phi$
neither $\texttt{tell}\left(\phi, \Xi^{i-1}\right)$ nor $\texttt{tell}\left(\neg\phi, \Xi^{i-1}\right)$	Ξ^i knows neither ϕ nor $\neg\phi$

Table 2: Epistemic states of a KB as a result of the operation `tell` over it.

It should be obvious by now that the two last rows of Table 2 are not *epistemic states* proper, as it is not possible to know both a sentence ϕ and its negation $\neg\phi$ (e.g. "Penguins fly" and "Penguins do not fly"), and if one knows neither a sentence ϕ nor its negation, then one simply has no knowledge with respect to ϕ , where "no knowledge" does not here mean ignorance, which was defined above by the justifier $\mathcal{J}_S^{\mathbf{U}} = (\mathcal{P}, \emptyset)$. If a KB is inconsistent, any sentence whatsoever can be entailed by it, a problem already mentioned above known as *ex falso quodlibet*, so that for an inconsistent KB Ξ, the number of *sentences* entailed by it is greater than the number of *facts* it does entail. If a KB is incomplete, then the reverse consequence is the case: It entails fewer facts than it should.

Knowledge Access & Share – By this it is meant here that the agent *queries* the KB with respect to what it knows. By replying to the query, the KB shares its knowledge with the KA (and this can further share it with other KAs). In its simplest form, `ask` is an operation that prompts the KB to a "Yes/No" answer, i.e.

$$\texttt{ask}\left(\phi, \Xi^i\right) \in \{yes, no\}$$

such that

$$\texttt{ask}\left(\phi, \Xi^i\right) = \begin{cases} yes & \text{if } \phi \in \Xi^i \\ no & \text{otherwise} \end{cases} .$$

By querying the KB, the KA is actually setting in motion the operation $\models$ by the KB's inference engine, so that we have:[27]

[26] Easier said than done, as negation in KBs by means of the operator $\neg$ is not a matter of fact. See below a possible solution to this problem.

[27] For simplicity, I am not separating the inference engine from the KB. This is possible in Prolog environments, in which the KB is in fact a program implementing a resolution calculus. See Augusto (2020a), Chapter 9, for details.

$$\texttt{ask}\left(\phi, \Xi^i\right) = \begin{cases} \textit{yes} & \text{if } \Xi^i \models \phi \\ \textit{no} & \text{otherwise} \end{cases}$$

Clearly, we have the equivalences $yes \equiv 1$ and $no \equiv 0$, so that in fact we can interchange $\{no, yes\}$ and $\{0, 1\} = W_2$. Figure 8 shows the complete deductive KS at hand.

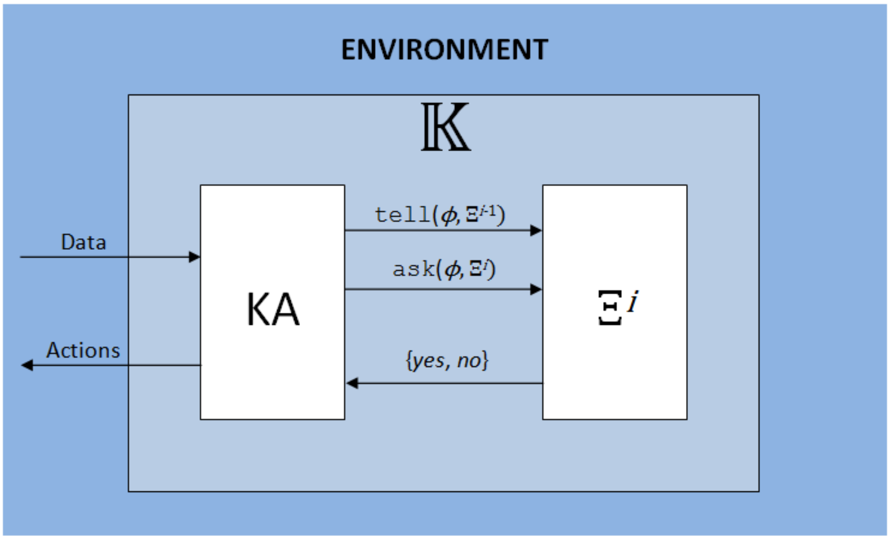

Figure 8: A deduction KS.

However, it is rarely the case that KBs are so "crisp" (to use jargon from fuzzy logics); as a matter of fact, a KB can hardly be expected to contain *all* the facts, and *only* the facts, with respect to some environment E. In particular, sometimes we are faced with sentences that are both true and false (a *contradiction* in classical logic; a *truth glut* in many-valued logics), and with sentences that appear to be neither true nor false (*truth gaps*). Moreover, the environment is in constant change, but KAs are not always aware of this, and may enter contradictory sentences at different times; also, it is more often than not the case that several KAs can use the operation `tell` on a single KB, so that one KA can add the sentence ϕ to the KB while another adds the sentence $\neg\phi$. Finally, KAs are typically not omniscient, so that relevant facts are often missing from KBs. The consequence of these scenarios is that KBs are often *de facto* incomplete and/or inconsistent. Hence, it might prove useful to augment the set $\{yes, no\}$ with $\{both, none\}$, so that the KB can reply in the cases that both $\phi \in \Xi^i$ and $\neg\phi \in \Xi^i$, or neither $\phi \in \Xi^i$ nor $\neg\phi \in \Xi^i$. In terms of truth-value set, we actually have $2^W = \{\emptyset, \{0\}, \{1\}, \{0, 1\}\}$, the power set of W_2.

Recall that the KB only knows explicitly what it is told by the KA, and it knows implicitly the facts that can be inferred by means of rules in the KB. Let us assume that the KA can tell the KB "negative" facts (e.g., "Penguins do not fly.").[28] Then, Table 3 (where superscripts are omitted in Ξ for simplicity) shows the different possible

[28] In fact, typically the KB cannot know "negative" facts, as the derivation of $\neg\phi$ is not without

$\models$	Reply	Epistemic set	Epistemic interpretation
$\Xi \models \phi$	*yes*	$\{K_\Xi\phi, \neg K_\Xi\neg\phi\}$	Ξ knows ϕ
$\Xi \not\models \phi$	*no*	$\{\neg K_\Xi\phi, K_\Xi\neg\phi\}$	Ξ does not know ϕ
both $\Xi \models \phi$ and $\Xi \not\models \phi$	*both*	$\{K_\Xi\phi, K_\Xi\neg\phi\}$	Ξ knows too much
neither $\Xi \models \phi$ nor $\Xi \not\models \phi$	*none*	$\{\neg K_\Xi\phi, \neg K_\Xi\neg\phi\}$	Ξ does not know enough

Table 3: Replies of a KB to `ask` (ϕ, Ξ^i) with corresponding epistemic states and epistemic interpretations.

cases with respect to the epistemic states of a KB, now represented by sets. (Recall that the inference engine is considered here as an internal component of the KB, so that the replies to `ask` are actually output by the engine.)

The pay-off of relinquishing the crispiness of classical logic is that a KB need not be discarded just because it is incomplete or inconsistent; as said, it is often the case that KBs are one or the other, or even both, and this solution allows us to reason epistemically within a KS in the presence of inconsistency and incompleteness.

A good example of this kind of reasoning, which additionally provides a means of reasoning about knowledge while avoiding the complications inherent to epistemic logic (e.g., Hocutt, 1972), is provided in Belnap (1977), where information and knowledge meet in the following interesting way: By replying "*yes*" to a query by the KA, the KB is *informing* the KA that *it has been told* that ϕ is a fact or that ϕ is true (abbr.: **t**), and by replying "*no*," the KB is telling that it has been told that $\neg\phi$ is a fact, or that ϕ is false (**f**); when replying "*both*" (**b**), the KB is informing that it has been told that both ϕ and $\neg\phi$ are facts or true, and "*none*" (**n**) is the reply given when the KB has been told nothing concerning the truth or falsity of ϕ. Then, and by resorting to the simple semantics of truth tables based on a valuation $val^4_{\mathcal{I}}$ for the superscript 4 denoting the truth-value set $W_4 = \{\mathbf{n}, \mathbf{f}, \mathbf{t}, \mathbf{b}\}$ (see Fig. 9), given a proposition such as $\chi \wedge \psi$ where $val^4_{\mathcal{I}}(\chi) = \mathbf{t}$ and $val^4_{\mathcal{I}}(\psi) = \mathbf{b}$, we have it that $val^4_{\mathcal{I}}(\chi \wedge \psi) = \mathbf{b}$.

	$\neg$
b	**b**
t	**f**
f	**t**
n	**n**

$\wedge$	**b**	**t**	**f**	**n**
b	**b**	**b**	**f**	**f**
t	**b**	**t**	**f**	**n**
f	**f**	**f**	**f**	**f**
n	**f**	**n**	**f**	**n**

$\vee$	**b**	**t**	**f**	**n**
b	**b**	**t**	**b**	**t**
t	**t**	**t**	**t**	**t**
f	**b**	**t**	**f**	**n**
n	**t**	**t**	**n**	**n**

Figure 9: Truth tables for Belnap's 4-valued logic.

Clearly, in face of this uncertainty, we are now in the context of information with respect to the KA. In other words, the replies output by a KB based on a logical language for whose interpretation there is a truth-value set W_k for $k \geq 3$ are *not* knowledge proper, but information in the sense that for each query concerning a sentence ϕ the KB can only provide the information that it has been given with

problems in a KB; the inference engine can only make the entailment $\Xi^i \not\models \phi$, which may then be interpreted as "$\Xi^i \models \neg\phi$," an interpretation known as *negation as failure*. Alternatively, $\Xi^i \not\models \phi$ just means that the KB *does not* know ϕ, so that ϕ is unknown, an interpretation known as *closed world assumption*. For central literature, I refer the reader to Reiter (1978) and Shepherdson (1984).

respect to ϕ. Because we have now $\mathcal{X} = \{x_i\}_{i=1}^{4}$ for the set $W_4 = \{\mathbf{n}, \mathbf{f}, \mathbf{t}, \mathbf{b}\}$ and $\Pr(\phi = x_i) = 1/4$, the entropy concerning ϕ is given by

$$H(\phi) = -\sum_{i=1}^{4} p_{\phi}(x_i) \log_2 p_{\phi}(x_i) = 2$$

where $p_{\phi}(x_i) \log_2 p_{\phi}(x_i) = -0.5$. Unsurprisingly, the entropy of a sentence that can be analyzed from the viewpoint of four equal probability outcomes is higher than that of a bivalent one. We have the following result: For two logical systems L_m and L_n with truth-value sets W_m and W_n, respectively, if $m \leq n$, then

$$H(\mathrm{L}_m) \leq H(\mathrm{L}_n)$$

if we consider m and n to be the equal probabilities that a sentence of these logics is valuated as $i \in W_{m,n}$.[29] Accordingly, for two KBs Ξ_m and Ξ_n in which a sentence ϕ can be interpreted over the truth sets W_m and W_n, respectively, we have for $m \leq n$:

$$H(\Xi_m) \leq H(\Xi_n).$$

However, it should be obvious that, when considered from the viewpoint of entropy, if $H(\Xi) > 0$ then Ξ is not a KB proper, but rather an information repository. Although only ideally does one have $H(\Xi) = 0$ for some KB Ξ, this remains as a general measure of entropy for a KB proper that works as a distinguishing feature between a KB and an information repository.

5 Conclusions and Further Work

The recent proliferation of the so-called knowledge communities brought about by IoK requires a general but precise notion of knowledge that can be used uniformly and consistently by all of them. IoK also brought with it a pernicious confusion of data, information, and knowledge (DIK) that needs to be effectively dissolved. These two problems are tackled in this original-research article.

In order to achieve the aimed-at generality the tools of mathematics are called for. Because I here adopt and adapt the Platonic definition of knowledge as justified true belief, whose foundation is the *proposition*, an admirable and intriguing merging of a mathematical and a linguistic objects, mathematical logic is my main tool. By means of this tool, a formal semantics for a logical language, coupled with an also formal notion of epistemic justification, is elaborated on that allows for a precise distinction between a *datum* $\phi = R(t_1, ..., t_k)$ and a fact with the same, or expressible in the same, form. It is often said that knowledge is *interpreted* data and I take here this in a wholly formal sense: ϕ is a *fact* if (i) for a relation symbol $R(t_1, ..., t_k) \in Rel_{\mathfrak{L}}$ and for every $t_i \in Ter$, $i = 0, 1, ..., k$, it is the case that there is an interpretation $\mathcal{I}$ and a corresponding valuation $val_{\mathcal{I}}$ such that

$$val_{\mathcal{I}}(R(t_1, ..., t_k)) = \Psi(R)(val_{\mathcal{I}}(t_1), ..., val_{\mathcal{I}}(t_k)) = 1$$

[29] More complex measures can be applied; for instance, Boričić (2017) uses countable partitions of a truth-value set over formulas of some logic L_m, obtaining $H(\mathrm{L}_2) \leq 2$ for classical logic and $H(\mathrm{L}_4) \leq 3.61$ for Belnap's 4-valued logic.

and hence there is a model $\mathcal{M}_{\phi} = (\phi, 1)$, and (ii) there is a justifier $\mathcal{J}_S = (\phi, J)$ such that if $J = \models$, then both the model and the justifier for ϕ coincide. If one can be assured that all the elements entered in a storage base are facts, then one is assured that one has a *knowledge base* (KB); furthermore, if rules–special complex facts–are added to this, one can obtain *new facts* by means of the operation $\models$, denoting adequate (i.e. both sound and complete) logical consequence, and of instantiation. *Facts*, atomic or complex, and *KBs* are thus the basic *knowledge structures*. (*Models* and *justifiers* are knowledge structures, too, but at the "metaknowledge" level.) When a KB is operated over by one or more *knowledge agents* (KAs) by means of operations that are in fact *knowledge processes* (KPs), we have a *knowledge system* (KS).

This formal exploration of (meta-)knowledge structures needs to be enlarged to yet other formal approaches, based, for instance, on order relations and their associated lattices, as well as on other formal structures such as closure/kernel operators and Galois connections. These have been shown to be naturally associated to the logical notion of logical consequence (e.g., Augusto, 2020d; Bonnay & Westerståhl, 2012; Caspard & Montjardet, 2003; Hardgree, 2005; Humberstone, 1996; Metcalfe et al., 2010), and they might provide further results with respect to knowledge–seen from a formal perspective–and its computation in KSs. The modal logics are explicitly based on order relations, but I only mentioned briefly one of the translations of their operators for necessity and possibility, to wit, epistemic logic, as this is more typically employed for reasoning about knowledge than for knowledge representation. However, the description logics, also based on the modal logics, are particularly relevant to the construction of a common representational ground for human and artificial KAs, so that they impact directly on the design of KBs. Moreover, they might provide new insights into the relations between probability/certainty and possibility/necessity as far as conditional facts (and associated rules) are concerned (see, e.g., Badie, 2020b).

Both data and facts are here postulated to be inert in themselves. From a mathematical perspective, we may consider that they are the case when $H(X) = 0$ for X some random variable that can in fact be a proposition when its truth value is unknown. $H(X)$ is in fact a quantification of uncertainty known as *entropy*, and whenever $H(X) > 0$ it can be said that one is in the presence of *information*. In particular, $H(X) > 0$ when some subject S, which is also a KA interacting with a KB within a KS, is called to act in a specific environment E upon X. In effect, every action entails some degree of uncertainty, with respect to either data or knowledge, the latter alone contributing to the maximum goal of an agent, to wit, *rational action*. Hence, and because KBs are often inconsistent and/or incomplete, it is also useful to be able to quantify the entropy that their underlying systems entail. Although some work has been done in this topic (e.g., Ellerman, 2018; Markechová et al., 2018), it has only recently started and its very foundations still need to be established.

As said, this work aims at generality; as such it is essentially formal, and thus does not discuss specific issues that arise from some of its contents. I next identify some of these issues and briefly elaborate on them.

As mentioned above, knowledge is naturally associated to cognition, so that to speak of a KA is often the same as speaking of a *cognitive agent*. Although a cognitive agent does not necessarily deal with knowledge (for instance, in perception, especially in early-stage perception), every KA is a cognitive agent, and this should be borne in mind in what follows. To begin with, the KA in a KS was here left unspecified as to whether it is a human or an artificial agent, but this is a central topic with repercussions in KBs. If the KA is a human, then it is reasonable to believe that the facts they enter in a KB are mental representations (beliefs), so that a KB is in fact a *base of mental representations*. This clearly determines the design of KBs, in particular at the knowledge-representation level. Recent work in KSs or components thereof emphasizes, or is actually based on, this aspect (e.g., Poldrack et al. 2011). If a KA is to be designed as a wholly artificial cognitive agent, typically called a *robot* or an *autonomous KA* (and an *autonomous KS*, by extension), this poses a plethora of specific issues in knowledge representation that need to be tackled (e.g., Paulus & Sun, 2019; Tenorth & Beetz, 2017; see Vernon, 2014, for a comprehensive discussion).

Agency issues also need addressing. From the narrower viewpoint of the Platonic notion of knowledge as justified true belief that plays a central role in this work, the question is posed of the *responsibility* of deciding that a datum is a fact and can be "told" to a KB as such. As elaborated on above, both truth and epistemic justification are largely independent of the KA–if seen from a formal perspective, as is the case here–but it is in effect the KA that "tells" a KB the facts it contains. How this formalism translates into "real" KAs, that is a topic for future work, and thus I leave the KA here as a black box. This work is expected to be interdisciplinary, involving at least epistemology and cognitive science, and mental representations will surely figure prominently in it, as it is my personal stance that both philosophy and cognitive science are largely–if not essentially–centered on the subject of mental representation (e.g., Augusto, 2006; 2013; 2014).

Additionally, there is the responsibility associated to the KA's action in the environment, which is supposed to be largely supported, if not motivated, by the KB. With respect to these issues, it appears that Dennet's (1987) humble *intentional stance* needs an enlarged revision taking into consideration factors such as individuality and normativity of agents (e.g., Barandiaran et al., 2009; Vermaas et al., 2013). Especially relevant is the conception of a notion of *values for autonomous KAs* and the problem of how to implement them as facts in a KB (e.g., Boissier et al., 2017; Dignum, 2017; Hooker & Kim, 2019). Humans, in particular, are autonomous KAs whose epistemic agency is however highly regimented by cultural and societal contexts, so that specific ethical issues related to knowledge, its creation and transmission, are posed from ethnological and anthropological viewpoints (see, e.g., Josephides, 2015, for a collection of discussions). No analysis of knowledge is complete without taking these aspects into consideration. In effect, human society at large is the largest, most complex KS of all.

Appendix: The Language $\mathfrak{L}$

A *formal language* is a structure of the type $\mathfrak{L} = (\mathcal{A}, G)$ where $\mathcal{A}$ is an alphabet, a finite or infinite set of symbols

$$\{\alpha_i\}_{i=1}^{k\leq\infty} = \bigcup \mathcal{A}_j$$

where the $\mathcal{A}_j$ are disjoint subsets of $\mathcal{A}$, and G is a grammar, a finite set of rules $\{\mathbf{r}_j\}_{j=1}^{m}$ governing the formation of the legal, or well-formed, constructs (strings) of

$\mathfrak{L}$.[30] I shall be working with the formal language $\mathfrak{L}$ such that $\mathcal{A}_{\mathfrak{L}}$, the alphabet of $\mathfrak{L}$, is defined by

$$\mathcal{A}_{\mathfrak{L}} = Rel \cup \underbrace{Var \cup Cons}_{Ter} \cup Op \cup Qtf$$

where *Rel* is an infinite set of relation symbols/strings of arity n (e.g., $Outcome(\overbrace{x, y}^{n=2})$), $Var = \{x, w, y, , ...\}$ is an infinite set of individual variables, *Cons* is an infinite set of symbols/strings denoting concrete entities (i.e. constants), $Op = \left\{\neg^1, \wedge^2, \vee^2, \rightarrow^2\right\}$ is a finite set of ^{m}-ary operators, and $Qtf = \{\forall, \exists\}$ is the set of the universal ($\forall$, read "for all") and existential ($\exists$, read "there is one") quantifiers over which the variables vary. Finally, a variable x or a constant c is a term $t \in Ter$. Because the elements of *Op* and *Qtf* are *logical constants*, the formal language $\mathfrak{L}$ is in fact a *logical language*. More specifically, $\mathfrak{L}$ is called a *first-order language* in logical jargon.[31]

The rules of $G_{\mathfrak{L}}$ concern the formation of expressions (strings of symbols) called *formulas*. Strings of symbols from $\mathcal{A}_{\mathfrak{L}}$ are said to be (well-formed) formulas if they are built according to these rules of $G_{\mathfrak{L}}$:

- An atom A is a relation symbol of arity $k \geq 0$ written as $R(t_1, ..., t_k)$ for $k \geq 1$ and simply r when $k = 0$. (A relation symbol P of arity 0 is called a *propositional variable*, and is commonly written in lower case as p.)
- An atom A is a formula ϕ.
- If ϕ is a formula, then $\psi = \neg\phi$ is also a formula.
- If ϕ, χ are formulas, then $\psi = \phi \wedge \chi$, $\psi = \phi \vee \chi$, $\psi = \phi \rightarrow \chi$ are also formulas.
- If ϕ is a formula and x a variable, then $\psi = \forall x\phi$, $\psi = \exists x\phi$ are also formulas.

If we denote a set of formulas by $\mathcal{P} \subseteq \mathfrak{L}$, then we can rewrite the above as $p \in \mathcal{P}$, $A \in \mathcal{P}$, $\phi \in \mathcal{P}$, $\neg\phi \in \mathcal{P}$, etc.

Acknowledgments

Farshad Badie and Walid S. Saba, who are both editors for the *Journal of Knowledge Structures & Systems* with a background in formal logic, commented on a draft of this paper. My thanks to them for their valuable feedback.

References

Artemov, S. & Fitting, M. (2019). *Justification logic: Reasoning with reasons*. Cambridge: Cambridge University Press.

Artemov, S. & Nogina, E. (2005). On epistemic logic with justification. *Proceedings of the 10th Conference on TARK*, 279-294.

[30]See Augusto (2020b) for a comprehensive discussion of formal languages from the syntactical viewpoint.

[31]$\mathfrak{L}$ may additionally have an infinite set of l-ary function symbols (e.g., $f(x)$) and the symbol for identity (=), but this extension dictates the undecidable character of this first-order language.

Augusto, L. M. (2005). *Who's afraid of idealism? Epistemological idealism from the Kantian and Nietzschean points of view.* Lanham, etc.: University Press of America.

Augusto, L. M. (2006). A little idealism is idealism enough: A study on idealism in Aristotle's epistemology. *Idealistic Studies*, *36*, 61-73.

Augusto, L.M. (2010). Unconscious knowledge: A survey. *Advances in Cognitive Psychology*, *6*, 116-141.

Augusto, L. M. (2011). Putting the horse before the cart: A pragmatist analysis of knowledge. *Trans/Form/Ação*, *34*(2), 135-152.

Augusto, L. M. (2013). Unconscious representations 1: Belying the traditional model of human cognition. *Axiomathes*, *23*, 645-663.

Augusto, L. M. (2014). Unconscious representations 2: Towards an integrated cognitive architecture. *Axiomathes*, *24*, 19-43.

Augusto, L. M. (2019). *Formal logic: Classical problems and proofs.* College Publications: London.

Augusto, L. M. (2020a). *Computational logic. Vol. 1: Classical deductive computing with classical logic.* 2nd ed. London: College Publications.

Augusto, L. M. (2020b). *Languages, machines, and classical computation.* 2nd ed. London: College Publications.

Augusto, L. M. (2020c). *Many-valued logics: A mathematical and computational introduction.* 2nd ed. College Publications: London.

Augusto, L. M. (2020d). *Logical consequences. Theory and applications: An introduction.* 2nd ed. College Publications: London.

Badie, F. (2020a). Logic and constructivism: A model of terminological knowledge. *Journal of Knowledge Structures & Systems*, *1*(1), 23-39.

Badie, F. (2020b). Towards contingent world descriptions in description logics. *Logic and Logical Philosophy*, *29*(1), 115-141.

Barandiaran, X. E., Di Paolo, E., & Rohde, M. (2009). Defining agency: Individuality, normativity, asymmetry, and spatio-temporality in action. *Adaptive Behavior*, *17*(5), 367-386.

Belnap, N. D. (1977). A useful four-valued logic. In J. M. Dunn & G. Epstein (eds.), *Modern uses of multiple-valued logic* (pp. 8-37). Dordrecht: Reidel.

Bernecker, S. & Pritchard, D. (2011). *The Routledge companion to epistemology.* New York: Routledge.

Boissier, O., Bonnet, G., Cointe, N., Mermet, B., Simon, G., Tessier, C., & de Swarte, T. (2017). Ethics and autonomous agents: Models for ethical autonomous agents. ANR ETHICAA – ANR-13-CORD-0006. Delivrable # 4.

Bonnay, D. & Westerståhl, D. (2012). Consequence mining: Constants versus consequence relations. *Journal of Philosophical Logic*, *41*, 671-709.

Boričić, M. (2017). A note on entropy of logic. *Yugoslav Journal of Operations Research*, *27*(3), 385-390.

Caspard, N. & Montjardet, B. (2003). The lattices of closure systems, closure operators, and implicational systems on a finite set: A survey. *Discrete Applied Mathematics*, *127*, 241-269.

Cooper, J. M. (ed.). (1997). *Plato: Complete works.* Indianapolis: Hackett.

Dalkir, K. (2005). *Knowledge management in theory and practice.* Amsterdam, etc.: Elsevier.

Dennet, D. (1987). *The intentional stance.* Cambridge, MA: MIT Press.

Dignum, V. (2017). Responsible autonomy. *IJCAI'17: Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 4698-4704.

Dzisah, J. & Etzkowitz, H. (2012). *The age of knowledge: The dynamics of universities, knowledge and society.* Leiden & Boston: Brill.

Ellerman, D. (2018). Logical entropy: Introduction to classical and quantum logical information theory. *Entropy*, *20*(9).

Frege, G. (1892). Über Sinn und Bedeutung. *Zeitschrift für Philosophie und Philosophische Kritik*, *100*, 25-50. (English trans.: On sense and reference. In P. Geach & M. Black (eds.), *Translations from the philosophical writings of Gottlob Frege* (pp. 56-78). 3rd ed. Oxford: Blackwell & Totowa, NJ: Rowman & Littlefield, 1980.)

Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, *23*(6), 121-123.

Goldman, A. (1967). A causal theory of knowing. *Journal of Philosophy*, *64*(12), 335-372.

Halpern, J. Y. (1995). Should knowledge entail belief? *Journal of Philosophical Logic*, *25*(5), 483-494.

Hardgree, G. M. (2005). Completeness and supervaluations. *Journal of Philosophical Logic*, *34*, 81-95.

Hendricks, V. F. (2006). *Mainstream and formal epistemology.* Cambridge, etc.: Cambridge University Press.

Hocutt, M. O. (1972). Is epistemic logic possible? *Notre Dame Journal of Formal Logic*, *13*(4), 433-453.

Hooker, J. & Kim, T. W. (2019). Truly autonomous machines are ethical. *AI Magazine*, *40*(4), 66-73.

Humberstone, L. (1996). Classes of valuations closed under operations Galois-dual to Boolean sentence connectives. *Publications of the Research Institute for Mathematical Sciences, Kyoto University*, *32*, 9-84.

Josephides, L. (ed.) (2015). *Knowledge and ethics in anthropology: Obligations and requirements*. London, etc.: Bloomsbury Academic.

Kidd, P. T. (ed.). (2007). *European visions for the knowledge age: A quest for new horizons in the information society.* Cheshire: Henbury.

Kripke, S. (2011). Nozick on knowledge. In *Philosophical troubles. Collected papers. Vol. I* (pp. 162-224). Oxford: Oxford University Press.

Levesque, H. J. & Brachman, R. J. (1987). Expressiveness and tractability in knowledge representation and reasoning. *Computational Intelligence*, *3*(1), 78-93.

Levesque, H. & Lakemeyer, G. (2000). *The logic of knowledge bases.* Cambridge, MA & London, UK: MIT Press.

Limbaugh, D., Landgrebe, J., Kasmier, D., Rudnicki, R., Llinas, J., & Smith, B. (2020). Ontology and cognitive outcomes. *Journal of Knowledge Structures & Systems*, *1*(1), 3-22.

Losee, R. M. (2017). Information theory for information science: Antecedents, philosophy, and applications. *Education for Information*, *33*, 23-35.

Markechová, D., Ebrahimzadeh, A., & Giski, Z. E. (2018). Logical entropy of dynamical systems. *Advances in Difference Equations.*

Metcalfe, G., Paoli, F., & Tsinakis, C. (2010). Ordered algebras and logic. In H. Hosni & F. Montagna (eds.), *Probability, uncertainty and rationality* (pp. 3-83). Pisa: Edizione della Normale.

Newell, A. (1982). The knowledge level. *Artificial Intelligence*, *18*, 87-127.

O'Donnell, J. J. (1992). *Augustine: Confessions.* Text and commentary. 3 vols. Oxford: Oxford University Press.

Paulus, S. & Sun, Y. (2019). A survey of knowledge representation in service robots. *Robotics and Autonomous Systems*, *118*, 13-30.

Poldrack, R. A., Kittur, A., Kalar, D. et al. (2011). The Cognitive Atlas: Toward a knowledge foundation for cognitive neuroscience. *Frontiers in Neuroinformatics*, *5*.

Ragsdell, G., West, D., & Wilby, J. (eds.). (2002). *Systems theory and practice in the knowledge age.* Boston, MA: Springer.

Reiter, R. (1978). On closed world data bases. In H. Gallaire & J. Minker (eds.), *Logic and data bases* (pp. 55-76). New York: Plenum.

Saba, W. S. (2020). Language and its commonsense: Where formal semantics went wrong, and where it can (and should) go. *Journal of Knowledge Structures & Systems*, *1*(1), 40-62.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*, 379-423, 623-656.

Shepherdson, J. C. (1984). Negation as failure: A comparison of Clark's completed data base and Reiter's closed world assumption. *Journal of Logic Programming*, *1*(1), 51-79.

Tarski, A. (1944). The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research*, *4*(3), 341-376.

Tenorth, M. & Beetz, M. (2017). Representations for robot knowledge in the KNOW-ROB framework. *Artificial Intelligence*, *247*, 151-169.

Vermaas, P. E., Carrara, M., Borgo, S., & Garbacz, P. (2013). The design stance and its artifacts. *Synthese*, *190*, 1131-1152.

Vernon, D. (2014). *Artificial cognitive systems: A primer.* Cambridge, MA & London, UK: The MIT Press.

Cite this paper as:

Augusto, L. M. (2020). Toward a general theory of knowledge. *Journal of Knowledge Structures & Systems, 1*(1), 63-97.