Axiomathes (2022) 32:459–475
https://doi.org/10.1007/s10516-021-09534-x

ORIGINAL PAPER

Meaning Relations, Syntax, and Understanding

Prakash Mondal[1]

Received: 1 August 2020 / Accepted: 6 January 2021 / Published online: 19 January 2021

Abstract
This paper revisits the conception of intelligence and understanding as embodied in the Turing Test. It argues that a simple system of meaning relations drawn from words/lexical items in a natural language and framed in terms of syntax-free relations in linguistic texts can help ground linguistic inferences in a manner that can be taken to be 'understanding' in a mechanized system. Understanding in this case is a matter of running through the relevant inferences meaning relations allow for, and some of these inferences are plain deductions and some can serve to act as abductions. Understanding in terms of meaning relations also supervenes on linguistic syntax because such understanding cannot be simply reduced to syntactic relations. The current approach to meaning and understanding thus shows that this is one way, if not the only way, of (re)framing Alan Turing's original insight into the nature of thinking in computing systems.

Keywords Meaning · Syntax · Understanding · Turing test · Meaning relations

1 Introduction

Linguistic meaning—whether lexical or phrasal or even discoursal–primarily rides on the syntax of natural language, mainly because it is syntax that primarily determines what form semantics would take given the arrangement of the lexical items. Thus, for instance, if the sentence 'We have to go to the stores to buy some goods' is structurally rearranged to '*To we have buy the to stores go goods to some', the syntactic structure is consequently altered with the meaning of the original sentence eliminated. In this way syntax provides the form upon which semantic structures are built. But this does not, of course, mean that syntax carries everything relevant to the (re)construction of meanings from the meaning-bearing elements of linguistic expressions. This is because certain contextually determined meanings cannot be fully extracted from syntactic relations, compositional or otherwise. This is

✉ Prakash Mondal
prakashmondal@la.iith.ac.in

1 Indian Institute of Technology Hyderabad, Kandi, Sangareddy, Telangana 502285, India

especially important for AI research, in that semantic processing in AI is heavily dependent on various lexical resources such as WordNet, FrameNet, VerbNet etc. and also on the construction of semantic representations. Thus, a need for shallow processing of linguistic meaning is more increasingly warranted. On the other hand, Searle's (1980) Chinese Room Argument has pointed to the insufficiency of syntax for semantics or meaning because syntax is not intrinsic to physical systems that compute (see also, Harnad 1989). Regardless of whether this is accurate or not[1] (see Baggini 2009; Schweizer 2012; Dennett 2013; Milkowski 2017), one thing is clear–syntactically compositional relations are not equivalent to semantic relations for computing. If this were so, the former would guarantee the implementation of the latter without any lacuna in anything indicative of understanding, and possibly vice versa. Notably, Turing (1950) in his well-known conception of the Turing Test, essentially a linguistic test, regarded communication with human beings as indicative of understanding. The test involves a computer and a human both hidden behind a screen or veil on the other side of which sits a human who is supposed to examine the linguistically framed responses from both the computer and the human in reply to questions posed by him/her. Both the computer and the human are certainly indistinguishable to the judge, since the judge does not know which response comes from whom. All that the judge does is check the verbal responses in response to his/her queries in order to tell the machine apart from the human. What is interesting here is that understanding is matter of communication through questions and answers. How a machine can actually mesh syntax with, or map it to, semantic structures is outside the purview of the test. Be that as it may, one significant point to note here is that the test does not simply equate understanding with the recovery of semantic structures from syntax or something of that sort, Rather, it leaves that issue open. Hence Descombes (2010) has even gone to the extent of saying that the Turing test is indicative of human gullibility rather than of understanding in machines. This is so because non-cognitive versions of computer programs can perhaps pass the Turing test by producing responses similar to those of humans (see Scheffel 2020). This particular issue has also been appraised by Warwick and Shah (2016), and they conclude that the test is not so much about checking human-type intelligence in machines as about testing some sort of communication ability in machines. Thus, there is reason to consider individuating the cognitive marks in heuristics and strategies that may be deployed by machines in a Turing test.

Since context, intentions and settings apart from syntax do contribute to linguistic meanings writ large, it seems worthwhile to consider how understanding can be construed in terms of meanings and operations them. The idea this paper advances is that meaning relations can be extracted from linguistic constructions in a manner that does not obey constraints of syntactic composition. Since natural language

[1] Exploring the validity or, for that matter, invalidity of the Turing Test is not the goal of this paper, and hence that question is outside the scope of the present discussion. We may note, however, that the Turing Test has targeted the understanding of natural language rather than some other cognitive task as the marker of thinking. This is what seems very relevant to the reformulation of the Turing Test in terms of the capacity to construct meaning relations to be discussed in the next section.

syntax is uniquely human (Berwick and Chomsky 2017; Terrace 2019), meaning relations need not be human, and hence there is no bar on their being attributed to machines (see Mondal 2017). On the present proposal, the lexical items of a language that can be regarded as the building blocks of substance-independent mental structures are indeed humanly manifest. The semantic or conceptual contents of lexical items are here considered to be part of the contents of substance-independent mental structures, given that the lexicon of a language contains and specifies disparate pieces of information that incorporate and integrate features of phonological, syntactic, semantic and possibly pragmatic properties (but see for a different view, Boeckx 2015). If so, there is no denying that the discrete structures constituting the minimal building blocks of mental structures are humanly realized in communities of language users (see also Terrace 2019). This underwrites the humanly realized character of the lexical items. Hence the minimal building blocks of substance-independent or organism-independent mental structures are humanly realized, but these mental structures in themselves are not. The larger structures composed from smaller elements (that is, lexical items) are thus non-unique, while lexical items are humanly unique in particular cases.[2]

Understanding in a mechanical system is a matter of running through meaning relations and executing relevant operations on them, thereby securing Turing's conviction that understanding in machines can be operationalized in a fashion that does not result in getting bogged down in the area of contention on the definition of thinking and understanding. The relevant conception of meaning relations is such that they can be shown to be constrained by a certain uniform logical structure of linguistic meanings across sentences and discourse contexts. It is this pattern of uniformity that can help, at least to some extent, gain a purchase on what we take to be understanding in a *minimalist* sense. The next section will show how to conceive of meaning relations and their relationship to syntax so that the notion of understanding can be framed in terms of meaning relations and relevant operations on them.

The paper is organized as follows. First, we outline the formulation of meaning relations and their relationship with syntactic composition. The nature of meaning relations is explicated by making reference to appropriate notions of syntax and compositionality. In Sect. 3 we describe the form of understanding in a mechanical system in terms of running inferences through meaning relations. In Sect. 4, we integrate the insights from the previous sections and offer certain apposite concluding remarks. This section also provides directions for future work.

[2] This contrasts with the familiar hypothesis that lexical items are the atomic elements that can be shared among humans, other animals and perhaps machines, whereas the structures built out of lexical items crossing the boundary between lexical items and other functional items (for example, prepositions, tense markers) are unique to humans (see Miyagawa et al. 2014). It needs to recognized that lexical items, when taken to be atomic elements as part of a formal system, are actually conceptually empty minimal items in their formal characterization. This allows lexical items to be shared among humans and other organisms, but keeps the structures built from them from being so shared.

2 Meaning Relations and Syntax

We shall first formulate the concept of meaning relations and its relationship with syntax in order to drive home a notion of understanding that is minimal enough to apply to cognitive systems in general. It may be reasonably believed that this notion will help gain a purchase on the notion of understanding the Turing Test has attempted to capture. In any case, it is important to emphasize that meanings from the meaning-bearing elements of linguistic expressions are exactly those that may be realized by patterns of conceptualizations (Langacker 1999; Jackendoff 2002), embodied simulations in the form of sensory-motor experiences having pragmatic effects (Colston 2019), schematic representations (Duffley 2020) and internal mental constructions for conjoining predicates (see Pietroski 2018). Significantly, since the meaning-bearing elements of linguistic expressions participating in the meaningful association of linguistic concepts have inferential roles to play, irrespective of the nature of syntactic contributions, the building blocks of any meaningful association of words as concepts bear inferential roles that are carried over to larger and larger associations (Brandom 1994; but see also Brandom 2007). Notice that it is also plausible to construe such inferential roles in terms of conceptual roles in a system of linguistic expressions such that these roles in meaning associations can be formally governed (see for related discussion Rapaport 2000, 2002). This will become clearer as the concepts of meaning relations and The Norm Reduction Condition are explicated later in this section.

That syntax does not carry everything relevant to the (re)construction of meanings from the meaning-bearing elements of linguistic expressions can be conveyed in a better way if we look at the issue of meaningful associations of linguistic concepts across boundaries, and beyond the constraints, of syntactic composition. The exploration of a not-so-typical kind of meaningful associations of linguistic concepts that may or may not supervene on syntax can be made, regardless of whether syntactic structures are thought of in terms of richly articulated functional structures and structural displacements (Chomsky 1995, 2000), or in terms of phrase structure grammars (Bresnan 2001; Jackendoff 2002), or even in terms of categorial combinations constraining possible permutations of syntactic categories (Steedman 2020). That is so because the common empirical content of syntax as may be apposite to certain conceptual associations of words may be invariant across varying conceptions of syntactic structures (see Johnson 2015). Thus, for example, the relation of a verb to its internal and external arguments or the relation of a preposition to its complement is invariant, no matter whether the relation is instantiated through a binary combinatorial operation Merge (Chomsky 1995) or through concatenation/unification (Bresnan 2001; Jackendoff 2002) or through functor-argument relations (Steedman 2019), or even through dependency relations (Tesniére 1959; Osborne 2019). It is to this *invariance* of syntactic relations we shall appeal in order to show how certain meaningful associations can be constructed beyond such syntactic relations. Let's consider the sentence 'The waiter over there looks at my friend in fear and I do not know why.' If this sentence is assigned a compositional syntactic representation

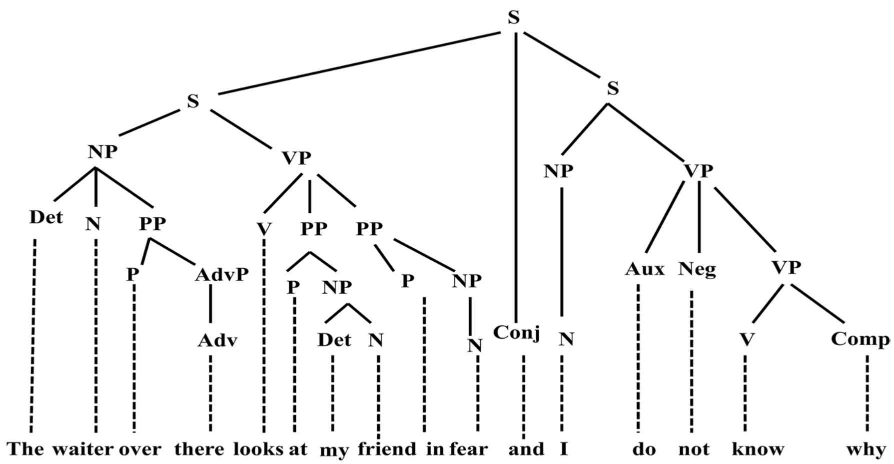

Fig. 1 The syntactic tree of the sentence 'the waiter over there looks at my friend in fear and I do not know why'

in terms of basic phrase structure, we can have the following representation by adopting a minimally simple form of syntactic representation (Culicover and Jackendoff 2005; Culicover 2013) which conforms to the basic constituent structure of sentences/clauses as can be mapped onto a functional structure in Lexical Functional Grammar (see Dalrymple, Kaplan and King 2016; Dalrymple and Findlay 2019) (Fig. 1).

It may observed that the tree diagram has certain constraints of syntactic composition on the grouping of words. For instance, 'the waiter over there' as a noun phrase (NP) is related to 'looks at my friend in fear' as a verb phrase (VP) in terms of direct composition but that is not the case between 'the waiter over there' and the verb phrase 'do not know why'. Likewise, the prepositional phrase 'over there' is in syntactic composition with the noun phrase 'the waiter' but not with the noun phrase 'my friend' or 'fear'. Additionally, the noun phrase 'the waiter' is not in syntactic composition with either 'my friend' or 'fear'. But how do we know this? The clear answer is that we know this precisely because syntactic compositions are explicitly displayed in the structuring of phrases and then their combination into a sentence. Semantic structures are assumed to ride on the way syntactic compositions function. That is why we know for sure that the waiter rather than the speaker looks at the speaker's friend since 'the waiter over there' but not 'I' is syntactically composed with the verb phrase 'looks at my friend in fear'. But certain otherwise valid combinations of linguistic concepts are explicitly ruled out by the constraints of syntactic composition here. For example, 'the waiter' and the prepositional phrase 'in fear' are conceptually related to each other but that cannot be captured in terms of syntactic composition. Likewise, 'the waiter' and 'my friend' may be conceptually connected to each other as a an experiencer-experienced pair. Similarly, 'the waiter' is also conceptually associated with 'at my friend' but this is not within the constraints of syntactic composition because 'at my friend' is composed with the verb directly and very indirectly with the noun phrase 'the waiter over there' as a whole (not specifically with 'the waiter'). Besides, there is a higher-order relation of a reason-situation pair between 'do not know' and the relation formed from 'the waiter' and 'in fear'. This too cannot ride on syntactic composition which exists between the verb 'know' and its sluiced clausal complement 'the waiter looks at my friend in fear' which in itself is syntactically composed. Thus, there remains a gap between syntactic composition-driven semantic structures/relations and conceptual associations that are not so grounded. It is the latter that is going to be useful in the characterization of the minimal building blocks of understanding. This is what we turn to now.

We may suppose that any linguistic meaning that can be constructed from various combinations of lexical items in*Lex* can be characterized in terms of some relation(s) drawn from among infinitely many relations defined on $\{Lex \ldots \bigcup \{Lex \bigcup R_1, \ldots, R_k\} \ldots\}$, given that $R_1, \ldots, R_k \subset Lex \times Lex \times \ldots \times Lex$ where k is an arbitrary number and the Cartesian product can be applied from one to an arbitrary number of times. Hence these infinitely many relations have the form $R_1, \ldots, R_k, R_{k+1}, \ldots, R_\infty$, where $R_{k+1}, \ldots, R_\infty$ are higher-order relations. Thus, whatever $R_1, \ldots, R_k$ are constructed on *Lex* form a union with *Lex* itself, and this can be repeated, if necessary. The tenets of this formulation have also been discussed in the context of semantic processing in AI systems (Mondal 2018). An example can make this quite clearer. For instance, if we want to construct a meaning of the phrase 'a beautifully painted picture of the river' from the lexicon of English, the lexical items 'a', 'beautifully', 'painted', 'picture', 'of', 'the', and 'river' from *Lex* can be related to each other in terms of meaning relations involving the relevant lexical items.[3] Thus, one meaning relation obtains between 'a' and 'picture'; a meaning relation between 'beautifully' and 'painted'; one *second-order* relation between the relation for 'beautifully' and 'painted' put together and 'picture', and another between 'a' and a meaning relation for 'beautifully painted picture' or even 'beautifully painted picture of the river' (because 'beautifully painted picture' or 'beautifully painted picture of the river' is already formed by a first-order meaning relation). It may be observed that the identification of meaning relations on *Lex* does not involve, or depend on, the way syntactic relations, and for that matter, semantic compositions, are defined on the hierarchy of a tree. However, some meaning relations may well correspond to those that can be defined on lexical items as determined by syntactic compositions. For example, one could imagine a meaning relation between 'a' and 'river' that does not form any syntactically defined constituent in the given phrase, or even comes from the way syntactic composition combines the given words. Likewise, 'painted' and 'river' can form a meaning relation of a property-object pair, which is not governed by the syntactic composition of the prepositional phrase (PP) 'of the river' and the noun head 'picture'.

[3] Note that this notion of relation is way different from the relations that can be constructed, as in model-theoretic syntax, for nodes in a tree (such as precedence or dominance relations) and for categories such as NP (Noun Phrase), VP (Verb Phrase), S (Sentence) etc. which are properties of nodes (see for details, Pullum (2013)). In fact, the relations $R_1, \ldots, R_k, R_{k+1}, \ldots, R_\infty$ encompass many *dimensions* (such as string adjacency, precedence, dominance and parent-of relations etc.) in terms of which linguistic constructions can be characterized.

Note that if the Montague-type notion of compositionality as a homomorphism between syntactic and semantic operations is applied (Dowty 1979; see also Hodges 2012), it would not be possible to derive such meaning relations from the homomorphism as these relations are outside the scope of the homomorphism concerned. This is more so if a direct way of composing such meaning relations from the available syntactic relations is warranted. That is because the notion of direct compositionality demands the formation of semantic values from the *immediate* syntactic units or constituents (Barker and Jacobson 2007). Thus, for example, 'painted' and 'river' can form a meaning relation, and yet its semantic value, however conceptualized, does not arise from any immediate syntactic constituent. However, it may now seem that a notion of compositionality that allows for the emergence of meanings that are not to be simply found in the parts of a complex expression would permit meaning relations. But this is not the case. Let's take a version of compositionality that fits this description. *Functional compositionality*, as opposed to *ontological compositionality* that warrants the addition of no new entities at the level of the whole, may admit of new meanings at the level of the whole that may not be straightforwardly found in the parts (Pelletier 2017). Notwithstanding the plausibility of this scenario, it is clear that even functional compositionality requires the syntactic combination of parts of an expression in standard ways–what is relaxed is the necessity to have all of the meaning of the whole derived *only* from the meanings of the parts, *but not* the necessity to syntactically combine the parts in a well-formed manner.

It is true that an*intrinsically* non-compositional approach to linguistic meaning can be specified in such a way that the elements of meanings have correlates in semantically composed expressions, as in models of linguistic meanings in which meanings are defined in terms of vectors defined over the features of words or combinations of words (see Turney and Pantel 2010; Clark 2015). Crucially, the present formulation of meaning relations, which may look similar to these models in having nothing to do with compositionality per se, is way distinct from these models of meaning. The reason is that meaning relations are underlyingly constrained on conceptual grounds, regardless of whether certain well-designated features of words match or not. Therefore, meaning relations are those relations that constitute conceptually viable elaborations or associations of linguistic contents expressed in words. In other words, meaning relations are those that are instantiated by conceptually constrained associations of a set of given words.[4] Crucially, the present formulation of meaning relations may dovetail well with the concept of partially saturated structures that as *constructions* blur the distinction between words and rules of grammar (Jackendoff 2002, 2020; Goldberg 2006, 2019). For instance, some constructions involve a mixture of fixed and free expressions (for example, the verb and

[4] Significantly, this account has differences from *frame semantics* (Fillmore 1976) or *scripts* (Schank and Abelson 1977) because the approach here is much more granular. For instance, a meaning relation between 'a' and 'river' in 'a beautifully painted picture of the river' cannot be easily captured by a syntactic unit or by a frame (or by a script), say, 'being a river' which will end up mapping the whole noun phrase and its grammatical function to the frame but 'a' and 'river' in are actually discontinuous frame elements (or script elements). Besides, frames may themselves participate in compositional operations in a construction through inheritance of information from daughter signs in local trees to mother signs (see Sag 2012).

the succeeding NP in 'My friend [$_V$ drank] [$_{NP}$ the whole night] [away]' are free variables with 'away' as the fixed expression). If so, it may appear that we can treat the linguistic expressions (words or phrases) that participate in a meaning relation as fixed expressions with any intervening material as a freely variable structure. Thus, 'a' and 'river' in 'a beautifully painted picture of the river' would be fixed expressions, while 'beautifully painted picture of the' would be a variable expression. But the problem with this is that 'beautifully painted picture of the' is not a valid syntactic structure or constituent, whereas the verb or the NP in 'My friend [$_V$ drank] [$_{NP}$ the whole night] [away]' is a licit syntactic category or unit. Even in resultative constructions such as 'She [$_V$ wiped] [$_{NP}$ the floor] [$_{AP}$ clean]', as discussed in (Jackendoff 2002), the verb or the NP or the adjective phrase (AP) is a legitimate syntactic category or unit. Likewise, the same problem would ensue from the supposition that 'a beautifully__ picture of the__' or even 'picture of the' is a valid free syntactic unit when 'painted' and 'river' form a meaning relation.[5] Conversely, treating 'painted' and 'river' as freely variable expressions with the intervening structure as a fixed expression would not also be of any help since there is no idiomatic expression in English that has or contains fragments such as 'a beautifully__ picture of the__' or even 'picture of the'. In any case, the differences between semantic values constructed from standard syntactic relations and meaning relations should be evident. That is also because syntactic relations in constructions (that is, relations of a syntactic argument to another) may well be dispensable,[6] although syntactic roles of arguments (such as subject, direct object, indirect object etc.) are essential as they may be guided by conceptual structures in constructions (Croft 2001). The following figure (Fig. 2) shows the difference between hierarchical structures of syntactic composition and meaning relations which are not necessarily compositional.

One may thus wonder if there are some reliable constraints that govern what meaning relations can be formed and help filter out illicit relations on this conception. One simple way of constructing meaning relations is to think of them as *filler-gap* relations in which one entity in R must be/contain a gap that is filled in by a filler. Thus, in $R_i = \{(x_1, y_1), \ldots, (x_n, y_n)\}$ x or y must be/contain a gap and the other would be a filler. In 'a beautifully painted picture of the river', which was our example described above, 'a' contains a gap (requiring something which is to be nominally specified) which is filled in by 'picture', whereas 'a' and 'beautifully' cannot form a meaning relation because even though 'a' contains a gap 'beautifully' is not a filler because it is itself a gap in requiring an adjectival predicate. The notion of gap here is thus more general than is recognized in linguistic theory, in that it may encompass arguments, complements, predicates including verbs or even relations formed through a filler-gap link

[5] It may also be observed that the freely variable structures even in idioms such as 'take (something) for granted' or 'cut (something) to the bone' (with the freely variable parts underlined) are not illicit syntactic categories or units (as they are all NPs), while the fixed expressions may be productively illicit in the language ('for granted' is possible, but 'for accepted' is probably not).

[6] Croft shows that a language such as Kilivila has a fixed order of NPs after a verb, and hence their relations to each other do not help decipher the syntactic roles of the arguments. Rather, the syntactic roles of the arguments are mapped onto their participant roles.

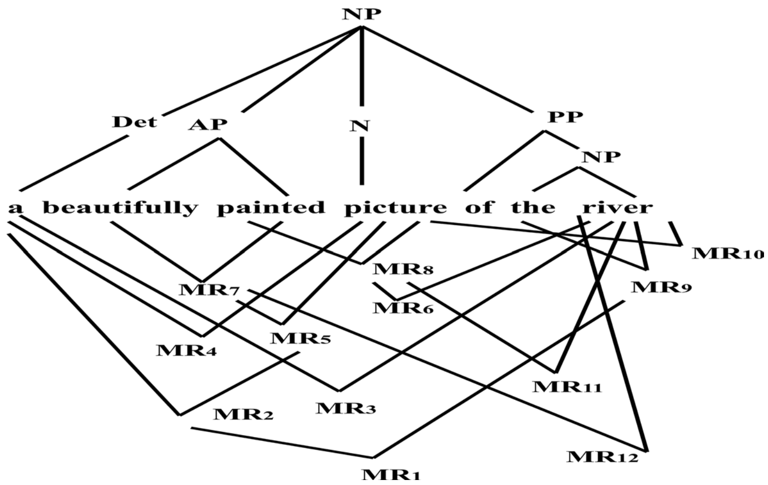

Fig. 2 The distinction between syntactic compositional structures and meaning relations ($MR_1 \ldots MR_{12}$): while MR_7, MR_8, MR_5, MR_2, MR_1, MR_9 can be derived in terms of relevant functions of syntactic composition, other MRs cannot

itself. Therefore, one rule of thumb is that any meaning relation must contain a gap and (at least) a filler. The asymmetry inherent in the relationship between the gap and the filler requires some amount of logical precision in description. Therefore, we can suppose that the construction of any arbitrary meaning relation R_i must result in the reduction of the cardinality (or norm) of the set characterizing the gap element or the filler element. We state this formally below.

2.1 *The Norm Reduction Condition for Meaning Relations*:

For any arbitrary meaning relation R_i if $(x, y) \in R_i$ where x and y are either the gap and the filler, or the filler and the gap, the cardinality or norm of the set characterizing x or y must be reduced when it is introduced into an *n-tuple* of R_i, that is, when it becomes part of $(x, y) \in R_i$ in the present case.

One way of providing a firm logical basis for the norm reduction condition above in order to secure the 'persistence' of information across meaning relations from the same array of lexical items in a given construction is to fix the exact logical role a gap plays in association with the filler in a certain meaning relation. The idea to be advanced here is that whenever a gap and a filler form a meaning relation the product of the formation must have a cardinality (or norm) reduced with respect to that of either the gap or the filler. Thus, for instance, 'picture' denotes a set of pictures and when 'a' and 'picture' form a meaning relation across a discontinuous distance of syntactic composition, the product as a set ('a' and 'picture' forming a meaning relation specifies a singleton set) has a cardinality reduced with respect to that of 'picture'. The words 'a' and ’beautifully’ cannot form a meaning relation, precisely because both ’a’ and ’beautifully’ are gaps, and even though ’beautifully’ specifies a set of ways or manners in which something can be beautiful, there is no linguistically available sense in which ’a’, if combined with ’beautifully’, forms a whole that reduces the cardinality of ’beautifully’. That is so because ’a’ does not specify a manner of being (beautiful) from among certain choices but does the job specification from among a set of entities. Likewise, the meaning relation (as a representation-object pair) constructed via the association of ’painted picture’ and ’river’ helps reduce the cardinality of the set characterizing ’painted picture’ (which is just a set of painted pictures) because the set of painted pictures *as representations* of rivers must be a smaller set than the set of all possible painted pictures. Equipped with the idea of meaning relations as the basis of minimal building blocks of understanding, we may now see how meaning relations help establish certain patterns of inferences that can be reliably identified with understanding in a computing system.

3 Understanding as Running Through Meaning Relations

We are now geared up to explore the nature of inferences as can be drawn from patterns of meaning relations. It is these inferences that permit the formation of the basic steps in any kind of reasoning. The formation of these step in reasoning may also be akin to the formation of minimal mental models (Johnson-Laird 2013), for inferences from patterns of meaning relations need not rest on representations of syntactic relations. It is of particular concern to stress that complex meaning relations formed out of simpler basic relations should retain the information the lower-level relations had. Taking the same initial example ‘a beautifully painted picture of the river’ we may illustrate this. If the meaning relation containing ’beautifully’ and ’painted’ (= way of being painted), and the relation involving ’a’ and ’picture’ (= specification of some picture) and then the one involving ’painted’ and ’picture’ (= quality or attribute of a picture) are all relations of the type R_0, the complex relation involving the basic relation for ’beautifully painted’ and ’picture’ would be an R_1 relation. The information contained in all relations of the type R_0 must be conserved, or would also persist, at the level of all those relations of type R_{0+n} that are made out of the exact R_0 and/or R_{0+n-1} type of relations that participated in the formation of the R_{0+n} type of relations. More significantly, a chain of deductions can be drawn up from the less embedded R_{0+n} type of relations to the more embedded relations of type R_0 or R_{0+n-1}. In this context, if a relation of type R_{0+n} is true, any relation that has participated at any lower level of embedding in the formation of that relation of type R_{0+n} is also true. More concretely, if the complex relation involving the basic relation for ’beautifully painted’ and ’picture’ is true, the meaning relation for ’beautifully painted’ must be true as well. This ensures that information percolates up to the less embedded relations in a *hierarchy* of meaning relations. This is shown below (Fig. 3).

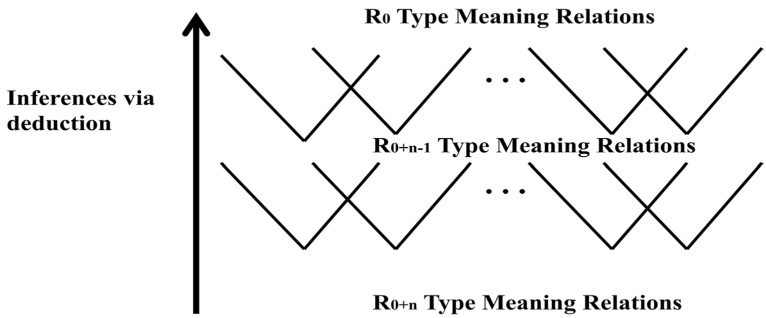

Fig. 3 The preservation of truth via deductive inferences from more complex to less complex meaning relations

Further, this comes with the consequence that the more lexical items participate in constructing meaning relations, the more ways of conserving information emerge through multiple levels of embedding of meaning relations. This also suggests that the lexicon is not simply populated with lexical items–rather, there are infinitely many conceptual associations that are built out of the lexicon but only some of them become gradually consolidated so as to allow for further combinations. Significantly, when such combinations arise they do so only to encapsulate and then transfer the information from less complex combinations to more complex combinations. This manifestation of inheritance is *nothing other than minimal understanding* in any kind of cognitive system that makes reference to, but actually goes way beyond, the specification of event structures, argument structures and qualia structures coded in lexical items and their typing operations over them (such as selection, accommodation and coercion), and that is because these operations are functionally compositional in the sense described above (see Pustejovsky 1995, 2012; Ramchand 2019). Thus, for example, the meaning relation formed with a meaning relation involving 'the' and 'waiter' on the one hand and another complex meaning relation for 'looks at my friend' on the other hand for our example sentence discussed in Sect. 2 will have a kind of event structure along with the relevant argument structure. This will be inherited at the level of more complex relations such as the one formed with a meaning relation involving 'the' and 'waiter' on the one hand and another complex meaning relation for 'looks at my friend in fear' on the other hand. Once qualia features of arguments along with the event structure information are inherited from lexical items through simpler relations to complex relations, the assimilation and interpretation of the basic linguistic properties of the relevant meaning emerges. This is exactly what underlies a minimal form of understanding which is characterized by the decoding of the conceptual properties of meaning from the basic linguistic properties. This is depicted below in Fig. 4.

Figure 4 shows that qualia features are inherited from R0-type relations all the way down to R4-type relations and also that the event-structural information will be inherited from R2-type relations to R4-type relations.

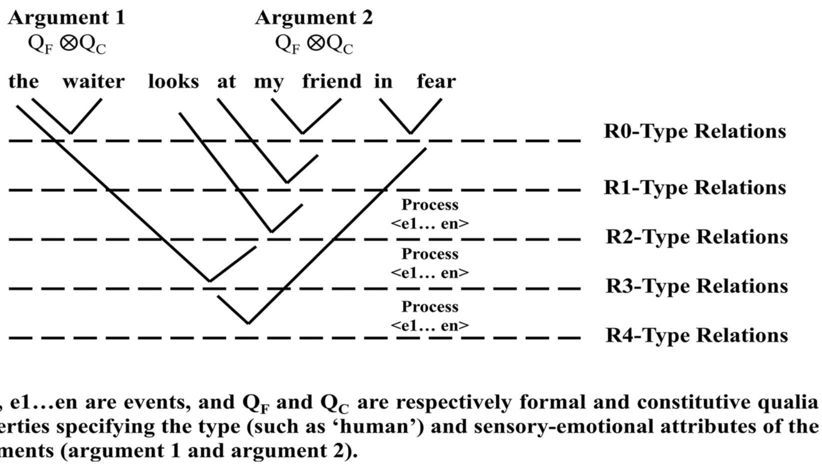

Fig. 4 The inheritance of event structures, argument structures and qualia structures across types of meaning relations

Thus, the construction of more complex meaning relations all the way from simpler or basic meaning relations and then running relevant inferences from more complex meaning relations towards those simpler meaning relations are the marks of understanding in a computing system that can construct simplex or basic meaning relations. The kind of understanding attributed to, and also equated with, the execution or implementation of inferences from more embedded or complex meaning relations to less complex ones can be the minimal building blocks of understanding in a mechanized system because one must be cautious in attributing simplistic mechanisms to more elaborate dialogs humans engage in (see Hofstadter 2016). It is also important to note that constructing and running through meaning relations, which is identified with understanding, does not have to appeal to *level-of* relations (see Boden 1988; Haugeland 2002), for meaning relations are neither a property of the level of programs nor a property of the level of the system concerned. Rather, meaning relations are facets of any fragment of a linguistic system that a cognitive system can (start to) evince when equipped with an inventory of lexical items (and possibly some concatenative rules of word combinations).

A caveat is in order here. It may not be clear what computing systems in themselves take to be the meaningful parts in their operations that yield certain outputs (from those parts) because the outputs may be produced based on non-linear patterns of learning not accessible to the modellers (see Nefdt 2020). Hence compositionality may not straightforwardly apply to the *processes* of computing systems constructing meaning relations. Nonetheless, the inference chains from complex meaning relations can preserve *process compositionality*. Process compositionality is exactly the specification of a compositional procedure in a system that is (process) compositional. While this holds, it remains true that meaning relations are not in themselves process compositional because some word associations may be (sometimes) unique or even idiosyncratic as chunks of concepts or pieces of knowledge.

Hence the term 'state compositionality' may apply to individual meaning relations because for state compositionality only identifying certain meaningful parts suffices. Taken in this sense, the construction of meaning relations simple or complex does not invite the same problem of compositionality because compositionality at the level of meaning relations does not, or need not, coincide with that at the level of machine operations. Hence this does not also hinge on the possibility of there being any abstraction relation between the hardware and the logic gate operations (Horsman et al. 2014; but see Szangolies 2020). That is because abstraction/representation relations at the level of machine operations in relation to the hardware and those that obtain for meaning relations may indeed differ. This ensures that meaning relations need not be tied to, or fully constrained by, aspects of a full-blown meaning-making system because the kind of understanding emerging from meaning relations is minimal at the least.

Be that as it may, the current way of framing understanding in a computing system has another benefit, in that it can also help arrive at certain abductions (Peirce 1998; see also Douven 2017) or 'guesses' that may have a probabilistic anchoring.[7] For instance, in the sentence described in Fig. 1 ('The waiter over there looks at my friend in fear and I do not know why.'), we may conceive of a complex or higher-order meaning relation between 'do not know' and 'my friend' (or even 'the waiter'). This relation would specify the knowledge state-agent pair whereby the person referred to by 'my friend' may also probably not know why the waiter looks at him/her in fear. Given the Norm Reduction Condition if applied across clause boundaries, the construction of this meaning relation is an obligatory result. But unlike other meaning relations (such as the one between 'do not know' and 'I' or the one between 'the waiter' and 'in fear'), the viability of this meaning relation does not ride on logical necessity even when the simpler meaning relations 'do not know' is built on are true. Clearly then, we have a case where it is *probable* that the person referred to by 'my friend' does not know whatever is being stated to be not known.

It appears vital to recognize that abductions are just the sort of inferences that may invite *meaning holism* (Quine 1953), and hence are thought to be non-computational in character (Fodor 2000; but see for a critique, Fuller 2019). But it is also worthwhile to note that arriving at abductions through meaning relations may have nothing to do with whether or not abductions are computationally manageable. So long as the construction of certain meaning relations is viable and plausible given conceivable situations, the plausibility can be cashed out in terms of the likelihood of certain situations. In fact, there is no computational way in which this can be prevented given the formulation of meaning relations we have at hand. Hence the question of whether machines can implement computations realizing abductions is inconsequential, insofar as constructing meaning relations for many natural language expressions invites abductions anyway. For instance, it has been observed that people form abductions for the conditional antecedent in certain cases of causal

[7] It is plausible that the probabilistic character of abductions may be compatible with Bayesian reasoning (see Lipton 2004; McCain and Poston 2014), but that is something which is outside the scope of this paper.

conditionals (such as ’If a dog has fleas, then it will scratch constantly’) (Politzer 2007). In such a case, the relevant meaning relation for the antecedent part would comprise a simple meaning relation with ’dog’ and ’fleas’, and ’if’. More linguistic evidence for abductions also comes from conditional-like constructions with ’(just) in case’, as in ’You should take an umbrella just in case it rains’, where the antecedent part ’just in case it rains’ is not actually a cause or condition for the action ’You should take an umbrella’. Rather, this is sort of a reason for the action or event described in the independent clause. Importantly, this reason is just a guess that it might rain, and hence it is abduction-like. The construction of such meaning relations is also another step towards abductions reliable or otherwise because abductive explanations do not guarantee truths anyway.

4 Conclusion

This paper shows how the existing AI thinking on understanding can be refined if the notion of thinking and/or understanding is (re)framed in terms of the construction and reconstruction of meaning relations. A formal analysis of the viability of this idea has been presented with its advantages for marking deductions and abductions as characteristic of (minimal) understanding. However, we have employed simplifying assumptions about the generation of linguistic meanings that have connections to language users’ hidden intentions, complex implicit aspects of contexts and settings. But hopefully, more refined implementations of this idea in computing systems can be considered so that Turing’s (1950) imitation game can be recast as a test of the production of more and more complex meaning relations. The production of meaning relations need not be anything like actual human thinking, as Turing cautioned against unnatural reasoning while arguing for *naturalness* in a machine’s responses as may be typically expected in humans but typicality does not imply identity.

> It might be urged that when playing the ‘imitation game’ the best strategy for the machine may possibly be something other than imitation of the behaviour of a man. This may be, but I think it is unlikely that there is any great effect of this kind. In any case there is no intention to investigate here the theory of the game, and it will be assumed that the best strategy is to try to provide answers that would naturally be given by a man (Turing 1950: 435).

We acknowledge the huge gap between actual understanding and such understanding as may be attributed to a system running through meaning relations. But it is this gap that will reveal what is missing.

Acknowledgements I’m grateful to the anonymous reviewer and the editor for their insightful comments on the article, and for their judicious assessment of the argumentation deployed in the article.

Authors’ Contribution Not applicable.

Funding Not applicable.

Availability of Data and Material Not applicable.

Code Availability Not applicable.

Compliance with Ethical Standards

Conflict of interest The author declares that have no conflict of interest.

References

Baggini J (2009) Painting the bigger picture. The Philosopher's Magazine 8:37–39
Barker C, Jacobson P (2007) Direct compositionality. Oxford University Press, New York
Berwick RC, Chomsky N (2017) Why only us: Language and evolution. MIT Press, Cambridge, MA
Boden M (1988) Computer models of the mind. Cambridge University Press, Cambridge
Boeckx C (2015) Elementary syntactic structures: Prospects of a feature-free syntax. Cambridge University Press, Cambridge
Brandom R (1994) Making it explicit. Harvard University Press, Cambridge, MA
Brandom R (2007) Inferentialism and some of its challenges. Res 74(3):651–676
Bresnan J (2001) Lexical functional syntax. Blackwell, Oxford
Chomsky N (1995) The minimalist program. MIT Press, Cambridge, MA
Chomsky N (2000) New horizons in the study of language and mind. Cambridge University Press, Cambridge
Clark S (2015) Vector space models of lexical meaning. In: Lappin S, Fox C (eds) Handbook of contemporary semantic theory. Blackwell, Oxford, pp 493–522
Colston HL (2019) How language makes meaning. Cambridge University Press, Cambridge
Croft W (2001) Radical construction grammar. Oxford University Press, New York
Culicover PW, Jackendoff R (2005) Simpler syntax. Oxford University Press, New York
Culicover PW (2013) Explaining syntax: Representations, structures, and computation. Oxford University Press, New York
Dalrymple M, Kaplan RM, King TH (2016) Economy of Expression as a principle of syntax. Journal of Language Modelling 3(2):377–412
Dalrymple M, Findlay JY (2019) Lexical functional grammar. In: Kertész A, Moravcsik E, Rákosi C (eds) Current approaches to syntax: A comparative handbook. Mouton de Gruyter, Berlin, pp 123–154
Dennett D (2013) Intuition pumps and other tools for thinking. W.W. Norton and Co, New York
Descombes V (2010) Mind's provisions: A critique of cognitivism. Princeton University Press, Princeton
Douven I (2017) What is inference to the best explanation? And why should we care? In: Poston T, McCain K (eds) Best explanations: New essays on inference to the best explanation. Oxford University Press, Oxford, pp 7–24
Dowty D (1979) Word meaning and Montague grammar: The semantics of verbs and times in generative semantics and in Montague's PTQ. Springer, Dordrecht
Duffley P (2020) Linguistic meaning meets linguistic form. Oxford University Press, New York
Fillmore CJ (1976) Frame semantics and the nature of language. Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech 280:20–32
Fodor J (2000) The mind doesn't work that way: The scope and limits of computational psychology. MIT Press, Cambridge, MA
Fuller TJ (2019) Cognitive architecture, holistic inference and Bayesian networks. Mind Mach 29:373–395
Goldberg A (2006) Constructions at work: The nature of generalization in language. Oxford University Press, New York
Goldberg A (2019) Explain me this: Creativity, competition, and the partial productivity of constructions. Princeton University Press, Princeton
Harnad S (1989) Minds, machines and Searle. J Exp Theor Artif Intell 1:5–25
Haugeland J (2002) Syntax, semantics, physics. In: Preston JM, Bishop MA (eds) Views into the Chinese room: New essays on Searle and artificial intelligence. Oxford University Press, New York, pp 379–392

Hodges W (2012) Formalizing the relationship between meaning and syntax. In: Werning M, Hinzen W, Machery E (eds) The Oxford handbook of compositionality. Oxford University Press, New York, pp 245–261
Hofstadter, D. R. (2016). Dull rigid human meets ace mechanical translator. In S. B. Cooper, A. Hodges (Eds.). *The once and future Turing*. Cambridge: Cambridge University Press.
Horsman C, Stepney S, Wagner RC, Kendon V (2014) When does a physical system compute? Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 470(2169):20140182
Jackendoff R (2002) Foundations of language: Brain, meaning, grammar, evolution. Oxford University Press, New York
Jackendoff R (2020) The texture of the lexicon. Oxford University Press, New York
Johnson K (2015) Notational variants and invariance in linguistics. Mind and Language 30(2):162–186
Johnson-Laird PN (2013) The mental models perspective. In: Reisberg D (ed) The Oxford handbook of cognitive psychology. Oxford University Press, New York, pp 650–667
Langacker R (1999) Grammar and conceptualization. Mouton de Gruyter, Berlin
Lipton P (2004) Inference to the best explanation, 2nd edn. Routledge, London
McCain K, Poston T (2014) Why explanatoriness is evidentially relevant. Thought 3:145–153
Milkowski M (2017) Why think that the brain is not a computer? APA Newsletter on Philosophy and Computers 16(2):22–28
Miyagawa S, Ojima S, Berwick RC, Okanoya K (2014) The integration hypothesis of human language evolution and the nature of contemporary languages. Frontiers in Psychology 5:1–6
Mondal P (2017) Natural language and possible minds: How language uncovers the cognitive landscape of nature. Brill, Leiden/Boston
Mondal P (2018) Lexicon, meaning relations, and semantic networks. In: Proceedings of the 2nd workshop on natural language for artificial intelligence (NL4AI 2018), Trento, Italy, pp 40−52
Nefdt RM (2020) A puzzle concerning compositionality in machines. Mind Mach 30:47–75
Osborne T (2019) Dependency grammar. In: Kertész A, Moravcsik E, Rákosi C (eds) Current approaches to syntax. Mouton de Gruyter, Berlin, pp 361–388
Peirce, C. S. (1998). *The essential Peirce: Selected philosophical writings,* Vol. 2 (1893–1913). Peirce Edition Project (Eds.). Bloomington: Indiana University Press.
Pelletier FJ (2017) Compositionality and concepts – A perspective from formal semantics and philosophy of language. In: Hampton JA, Winter Y (eds) Language, cognition, and mind, vol 3. Compositionality and concepts in linguistics and psychology. Springer, New York, pp 31–94
Pietroski P (2018) Conjoining meanings: Semantics without truth values. Oxford University Press, New York
Politzer G (2007) Reasoning with conditionals. Topoi 26:79–95
Pullum GK (2013) The central question in comparative syntactic metatheory. Mind and Language 28(4):492–521
Pustejovsky J (1995) The generative lexicon. MIT Press, Cambridge, MA
Pustejovsky J (2012) Type theory and lexical decomposition. In: Pustejovsky J et al (eds) Advances in generative lexicon theory. Springer, Berlin, pp 9–38
Quine W (1953) From a logical point of view. Harvard University Press, Cambridge, MA
Ramchand, G. (2019). Event structure and verbal decomposition. In R. Truswell (Ed.). *The Oxford Handbook of Event Structure*. New York: Oxford University Press.
Rapaport WJ (2000) How to pass a Turing test: Syntactic semantics, natural-language understanding, and first-person cognition. J Logic Lang Inform 9:467–490
Rapaport WJ (2002) Holism, conceptual-role semantics, and syntactic semantics. Minds and Machine 12:3–59
Sag I (2012) Sign-based construction grammar. In: Boas HC, Sag IA (eds) Sign-based construction grammar. CSLI Publications, Stanford, pp 69–202
Schank RC, Abelson RP (1977) Scripts, plans, goals, and understanding. Lawrence Erlbaum, Hillsdale
Scheffel J (2020) On the solvability of the mind–body problem. Axiomathes 30:289–312
Schweizer P (2012) The externalist foundations of a truly total Turing Test. Minds and Machine 22:191–212
Searle J (1980) Minds, brains and programs. Behavioral and Brain Sciences 3:417–457
Steedman M (2019) Combinatory categorial grammar. In: Kertész A, Moravcsik E, Rákosi C (eds) Current approaches to syntax. Mouton de Gruyter, Berlin, pp 389–420
Steedman M (2020) A formal universal of natural language grammar. Language 96:1–43

Szangolies J (2020) The abstraction/representation account of computation and subjective experience. Minds and Machine 30:259–299

Terrace, H. S. (2019). *Why chimpanzees can't learn language and only humans can*. New York: Columbia University Press.

Tesniére L (1959) Eléments de syntaxe structurale. Klincksieck, Paris

Turing A (1950) Computing machinery and intelligence. Mind 59:433–460

Turney PD, Pantel P (2010) From frequency to meaning: Vector space models of semantics. Journal of Artificial Intelligence Research 37:141–188

Warwick K, Shah H (2016) Passing the Turing test does not mean the end of humanity. Cognitive Computation 8:409–419