An Active Structure-Based Artificial General Intelligence

Jim Brander * and Ryan Sigmundson
Interactive Engineering, Sydney, NSW, Australia.

* Corresponding author. Email: jimbrander@gmail.com (J.B.)
Manuscript submitted March 30, 2023; accepted April 19, 2023; published April 25, 2023.
doi: 10.18178/JAAI.2023.1.1.1-19

Abstract: In this article, we introduce an artificial intelligence-based approach to numerical and logical analysis that relies upon a network built up of semantic structures. Our design utilizes an undirected, fully transparent architecture with a capacity for self-extension. It is designed to process information and reason using natural language as its substrate. The network, which we refer to as Active Structure, consists of a semantic network derived from tens of thousands of definitions extracted from a human-readable dictionary. Extensive curation of the dictionary entries was required to translate them into a machine-interpretable form. We outline our solutions to issues that were commonly encountered during the curation process such as circularity, unfounded hierarchical organization, and inadequate detail. We go on to discuss possible applications of Active Structure, including an in-depth example of how it could be employed in the detection of money laundering. An overarching goal of our project is to design a machine intelligence that can overcome the Four Pieces Limit—an upper limit on the capacity of the human mind for conscious processing. We view surpassing this limitation as a crucial step toward the generation of new solutions to a wide variety of problems.

Keywords: Semantic structure, human-readable dictionary, undirected active network, self-extensible, dictionary curation, four pieces limit, anti-money laundering.

1. Introduction

Language is something that we use effortlessly throughout the day as we assemble and interpret sentences. For the vast majority of language users, comprehension occurs near-instantaneously without deliberate effort. That being said, the processes involved in language comprehension are numerous and complex. Setting aside obvious problems such as disentangling the influence of context on meaning [1, 2] or making sense of figurative language [3], even determining the syntactic role a given word plays in a sentence can prove to be a huge obstacle [4, 5]. The closer you look at it, the more the problems involved in language interpretation multiply.

For instance, the problem of polysemy, whereby a single word can have many different meanings [6, 7]. How does the mind go about determining what particular meaning of a word is intended in any given instance? The simple answer is that it uses context to resolve such ambiguities, but the myriad ways that context may interact with a word to resolve the intended meaning is anything but simple. Nevertheless, the human mind generally succeeds in carrying out the process of context-based lexical disambiguation within 200ms during spoken language comprehension [8, 9].

Attempts to design text interpretation machines capable of resolving lexical ambiguities have thus far failed to reach human levels of accuracy [10], in part because of the opacity of the system that they're trying to replicate. Even if such a machine did succeed at achieving a human level of disambiguation accuracy, the human mind itself is far from perfect, and misunderstandings do occur. Furthermore, lexical ambiguity resolution has been shown to be affected by age [11, 12] as well as working memory load [13]. It is therefore a system that could be improved upon, and doing so may avoid costly misinterpretations in the future.

Though an impressive language processor in comparison to machines, the human mind falls short in other capacities. There are severe limitations to the amount of information that can be held within working memory for conscious processing at any given time–a Four Pieces Limit, to be exact [14]. This is a severe limitation, and applies just as much to the people developing the solution we will describe, as to the problems the solution is meant to handle. More than four pieces and people unconsciously begin to chunk the problem to stay within the limit, making parts of the problem static, which affects the generality and longevity of any solution they find. There had been no point talking about the Four Pieces Limit until now, when there appears to be a way around it. Our mental load suffers while reading complex technical specifications, bureaucratic, and legal documentation, which can all too readily lead to mistakes. Misinterpreting documents of this sort can yield hefty consequences that could be readily avoided through use of a language-competent machine capable of overcoming the limited processing capacities of the human mind.

Given the complexity of even basic sentence parsing, the implication of such a limitation is that the majority of processing must take place in the unconscious mind. This fact means that in order to replicate the processes going on in our own brain we need to work backwards from our outputs. The field of psycholinguistics has been engrossed in this undertaking since its conception as an area of study [15], which to this day still heavily influences some artificial intelligence approaches to language processing [16–18], though the bulk of researchers in the area have moved on to deep learning (DL) approaches [19, 20]. Though DL itself is based on the utilization of artificial neural networks [20, 21], an idea closely linked to cognitive science, such networks bear no more than surface similarity to their counterpart in humans. Their focus is more on achieving functional outputs than on learning and processing in directly comparable ways to the human mind. Though successful in many constrained problem spaces, DL has some notable shortcomings, such as its difficulty in extrapolating beyond the bounds of its training space [21].

A division within the field of artificial intelligence can be drawn between narrow AI and artificial general intelligence (AGI). While narrow AI focuses on addressing a specific problem or domain, AGI has the capacity for broad generalization, allowing such systems to carry out a wide range of tasks in many different contexts [22].

A key feature of AGI is that it must be able to continue gathering and linking together knowledge in new ways–it is through this process that it has the potential to produce solutions that its creators may not even have conceived of. An obstacle in achieving this is determining how the machine will go about gathering its own knowledge–this brings us back round to language processing. The vast majority of human knowledge is encoded in language [23] and available in some form online. A machine capable of compiling and utilizing such resources would be capable of tackling problems far beyond its original parameters. In our view, the development of a machine capable of comprehending and flexibly making use of language is a key step toward competent AGI.

Throughout our paper, we therefore propose an approach to language processing that represents meaning through construction of initially undirected relational structures between words, groups of words, and groups of groups – all the way to words that act as existential operators with the potential to extinguish the existence of an entire document. We begin by suggesting some applications that a language-based AGI machine could potentially address. Next, we outline a variety of the details surrounding our modelling approach, including ways of grouping words and the modelling of words with multiple definitions, as well as their flexible and accurate selection from a potentially large definition set. We call our system architecture

Active Structure, which at its core is an undirected, self-phasing and self-extensible network. Finally, we use money laundering detection as an illustrative example to demonstrate how our approaches could be of use in a real-world application.

1.1. Planned Applications

The system that we will describe is well-suited to problems that are complex and/or multidisciplinary, and where data is peripheral or at least not central to the decision-making, is unavailable for a new problem, or is out of date for a problem that is following an underlying phenomenon that can only be inferred.

We begin by illustrating some possible applications off the AGI we are envisioning:

- **Complex Specifications** – a specification for a complex piece of technology can run to a thousand pages, have a glossary and terms defined inline, and be full of bullets referencing other bullets within clauses. In other words, a description that is hard to understand, even by experts. It will typically involve multiple areas of expertise – for a plane, that might mean wings, fuselage, engines, control surfaces, and avionics. If the technology is "new", there is no body of text to support it.

- **Complex Legislation** – similar to a complex specification, but with an important increase in difficulty. The people who pass it into law are not specialists, and are unlikely to see the legislation in any holistic way. Something that does see the legislation holistically, and can rapidly (within hours) show what happens if one or more aspect is changed, should be useful.

- **Strategic Planning** – an obvious example is the changeover from Internal Combustion Engine (ICE) vehicles to electric vehicles (EVs). The timeframe for the transition is long and there are many facets to be considered throughout the shift. For example, there is resistance from the public, a legislated timeframe for the transition after which ICE cars will not be sold, scarcity of the minerals needed, parallel production of ICE and EVs, and an eventual large reduction in the auto manufacturers' workforce. Furthermore, alternative technologies such as solid-state batteries and hydrogen fuel cells may emerge during the transition, so planning needs to be dynamic and easily reorganized over a long period of time.

- **Content Mediation** – Social media content is often spun around some new and newsworthy event, meaning that mediation cannot rely on a large block of text as input data to make sense of the social media post. Instead, it is necessary to analyze the post based on the meanings of the words, that is, at a micro level, with some foreknowledge about political trends and sarcasm.

- **Economic Modelling** – "Economics is less amenable than physics to definitive mathematical analysis because it is about people, whose responses to similar phenomena change over time. We build models in our minds or computers that fit observed reality at one point in time, and reality changes. Then we have to think harder about what's going on." – Ross Garnaut. The recent indecision in the US over whether the surge in inflation would be transitory or long-lasting illustrated the fragility and poor predictive power of economic models that do not merge psychology with more mathematical forms of modelling.

- **Climate Change** – a problem which is large and looming. The current data does not reflect the scale of the problem, and instead one or more climate models need to be used to estimate the scope of the problem over many years. There are many cross-specialties involved – atmospheric modelling, oceanographic modelling, agriculture in a changing climate, natural disasters.

A problem related to those above, and whose solution will be used by them, is Semantic Searching. By that we mean the ability to create a small piece of text describing precisely what you are looking for (by having the specific meanings identified) and have it search through potentially millions of Google hits using the meanings of words.

A problem that demonstrates the temporary hack of the human cognitive system in response to a textual command is Lane Following, during which we flexibly update our mental rules with regard to which lines to drive within in response to a road sign with instructions to overwrite our default behavior [24]. Modelled problems will require a similar facility when exploring different strategies – that is, a high-level instruction changing, extending and rebuilding the underlying structures, which implies everything is made of the same stuff and is navigable, not black boxes.

People can describe much more complex situations in English than in a contrived language, because English has had to include all the ways of describing the natural world, the world of human artifacts, and the mental world. Its great strength is that most of the processing of it is done unconsciously.

1.2. An Origin Story

We began work on a system combining arithmetic, algebra and propositional logic in 1982. The algebraic statements were equations, whose structure allowed propagation of values in any direction, as one would see in a simulation – when a plane is on the ground, the fuselage supports the wings, and when it is in the air, the wings support the fuselage. Switching could also occur through EQUALS and assignment operators, implication (IF THEN), Modus Tollens, or logical equivalence (IFF). A ground (similar to the surface of the whiteboard shown in Fig. 1 provided logical truth – an otherwise unconnected EQUALS operator would be connected to the logical ground. The system could employ Constraint Reasoning, and then recover its original state.

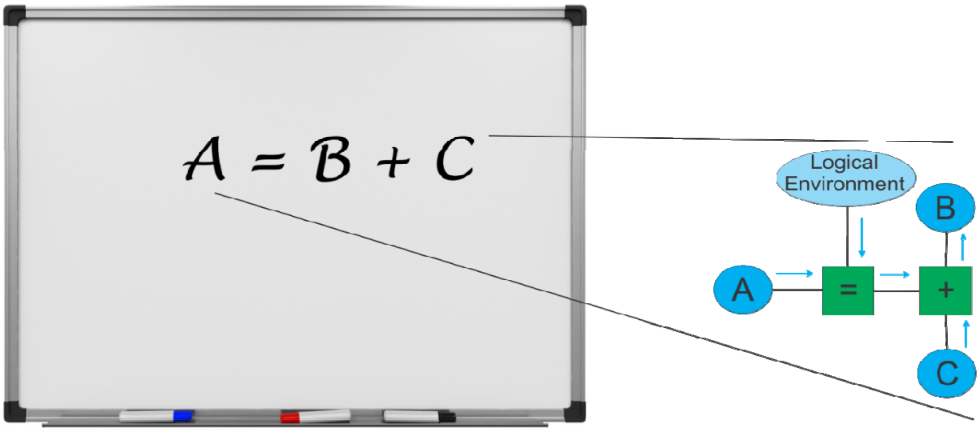

Fig. 1. Depiction of the logical structure A = B + C.

Marcus, in The Algebraic Mind [25], suggests that language manipulation can be an analogue of algebraic rearrangement – we don't see it that way, as language is vastly more complicated than anything that algebra can do. The simplicity of algebra can be seen in a quadratic equation, where X appears twice, but it is the same X. In language, the phrase "Fred was a lovely baby back in Indiana, but look at him now, fifty years later, with pain etched into his face" refers to the same Fred, but somehow not the same. In an algebraic network, multiple appearances of the variable in an equation can be linked so the variable only sees one connection. In a sentence, where the same object may be at different points on its timeline in different locations, it is not so easy. An undirected network–that is, a network in which the direction of information flow is not fixed–does not need algebraic rearrangement, and uses the states and values in the structure to work out which way information flows (in the same way that a pilot is not rearranging equations for take-off and landing – they have built a smooth simulation of the physical processes in their mind). As time passed, we realized that a similar structure could be used for language processing. We added existential logical connections to the non-numeric operators to mirror the simple way that English handles existential logic and the combination of temporal, propositional and existential logic (after the rain, the river is deep and he can't swim), and also the ability to switch blocks of text on and off (if gender is male, Section 9 is void). Constraint Reasoning turned into the ability to create new or amended structure, evaluate or use it, and dissolve it again (in the way of the Lane Following problem).

We designed a machine that could read health insurance policies and build an undirected structure to be used to answer questions about the contents of the policies (Is Botox covered?), where states and values could flow in any direction. Several problems emerged:

1) Words can have multiple meanings – in medicine, cervical vertebrae and cervical cancer. A colonoscopy cost $4,000, a misdiagnosis cost him his life. If there are only a few examples, some improvisations will work, but that is obviously not suitable for important or life-critical systems.

2) Where a word can be a noun or a verb, it can be very difficult to determine which it is, grammar being a very poor guide in anything but very simple text.

3) Unravelling preposition chains is almost impossible without access to meaning. And the meanings can be confusing – a group of people, a member of a group.

4) Sophisticated clumping is necessary. "An imaginary flat surface of unlimited extent" – the flat surface of unlimited extent is the thing that is imaginary.

We undertook to load definitions from a well-known dictionary into the machine, with the problem that the dictionary expects the reader to be sentient, while a machine with no initial definitions to guide it is the opposite. We have loaded 85,000 definitions to support a 36,000 word vocabulary, with definitions per word ranging from one to eighty two ("run"). Not all would be built initially, but we needed to be sure there were no insuperable problems with what we were doing – the near insuperable problems are with seemingly simple multi-use words, like "how". We are using Anti-Money Laundering (AML) as the Proof of Concept project – it would need about 18,000 definitions to get it off the ground.

2. Model Description

2.1. Relations as Objects

We will start to describe the proposed solution by talking about relations as objects, and using all the meanings of a word. Firstly, objects are more like relations than we like to think about. A car is a pile of spare parts without the assembly relation that holds it together. A person is a mass of processes that keep billions of cells functioning. Relations are represented as objects which can connect to other objects (which can be relations).

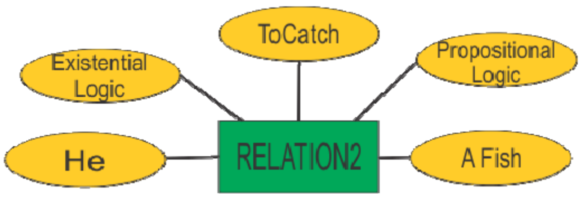

Fig. 2. Representation of the phrase "he caught a fish" as a relational object.

The particular RELATION operator in Fig. 2 has a subject and object (the 2 following the RELATION), and has propositional logic and existential logic connections, making it easy to say "he caught a fish", "he can't catch a fish", etc. Existential logic is used to control about 5% of statements internally, and all statements externally, so we can't just leave it out. The two logical connections can be used reversibly – "Did he catch a fish?", "Can he catch a fish?" – a search is instigated out from the operator on the appropriate connection to find out. Time is handled internally in the operator (it may be an instantaneous time for an instantaneous relation, or a time range, or connections to other relations which control the starting and stopping times), tense is handled by inheritance. The links can carry information in the form of states, numbers, objects, in any direction dictated by the operators, and operators can make or break connections, or create or destroy structure (up to about 40 million network elements).

Some RELATION operators produce a logical output:

"He knew there would be trouble and he would be blamed."

The above statement either needs an implied “for it” to be added, or the AND made a sequential AND, so the link between the clauses is made clear - the machine spends a lot of its time putting back the implied text that can be left out for people with an unconscious understanding of language rules.

The “knew” receives a logical true from the environment, and passes on True through a logical AND (the AND propagates logical states) to the subordinate clauses. A structural AND (it holds things together, but does nothing else) is used to connect the clauses to the object connection of the ToKnow relation (Fig. 3).

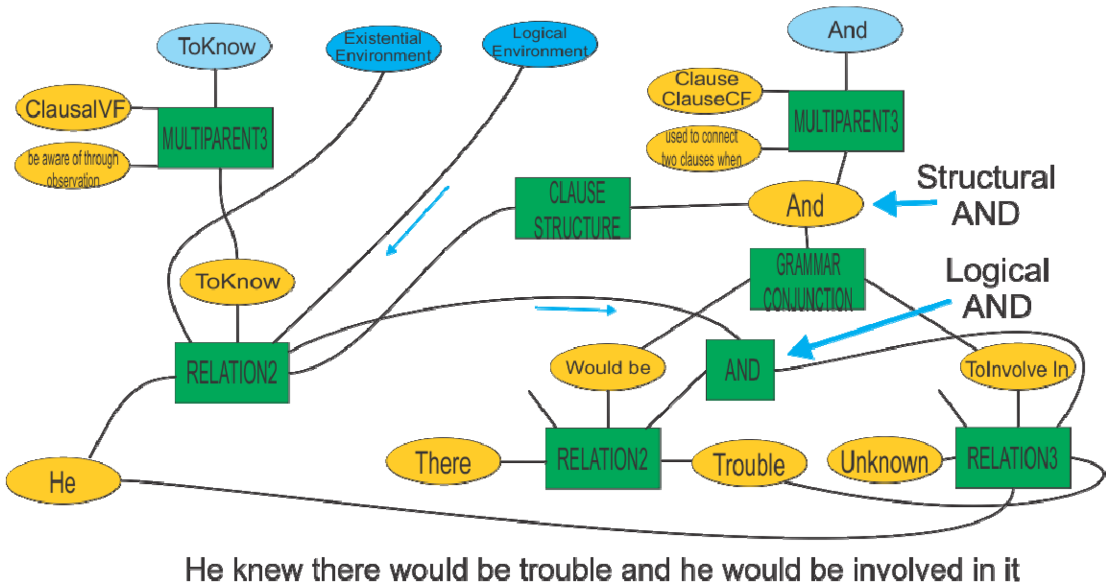

Fig. 3. The figure serves to illustrate the different usages of the Structural AND, which serves a grammatical function in connecting clauses, and the Logical AND, which establishes the existential equivalence of relations in the behind-the-scenes representational structure.

The logical AND isn’t really a logical AND as an electronics expert would understand it – it is functioning in reverse, as a splitter, but it can energize its normal inputs as outputs if it has a True on its nominal output. “OR” is a different matter entirely – it can only assume at least one of its inputs is True if it has a True on its nominal output, which in human language it handles easily.

There is ongoing discussion about whether humans have an innate ability to manipulate symbols, or whether it is learned [25, 26]. What is clear is that humans are very good at grouping things, rapidly assimilate logic, and can abstract things easily. As an example, take the definition of “by” - after a passive verb, where the object of the preposition can function as the subject of the verb. The child doesn’t ask how the word “verb”, acting as a noun, can be used to build a structure acting on verbs. In their mind, it is “obvious”. A machine without unconscious language processing has to be protected by the tagger detecting grammatical objects acting as nouns, make them honorary nouns, parse and build the structure, then convert them back to the objects they represent so the structure will be used as intended. But coming back to symbols and algebra – we think this approach is based on a misunderstanding. A separate “brain” does not manipulate words as symbols. The words represent pieces of machinery, which are duly connected according to the rules of the language, and operate as they should. There is no need for an external, complicated, and preconfigured device to manipulate them. An analogy to the process of language comprehension can be made with finding all the pieces of a mechanical clock on your desk, and the instructions to assemble them. You do so, and the clock runs. Lego is an even more apt analogy – the same pieces can be used to build all sorts of different physical things, but words are not limited to physical things, and the machine has some 36,000 different colored and shaped words to play with. There is an analogy with a jigsaw puzzle, except here the words change their color and shape based on the words around them.

2.2. WordGroups

The machine also has to handle groups of words that can be operated on as a unit. We call them Wordgroups, and currently have about 15,000 of them. A simple example would be “ambient temperature”. A more complex example is “at large”, which has a specific meaning (“a murderer at large”), or it may be used in another way, without the special meaning, as in “a DJ at large gatherings”. The machine needs to determine whether the special meaning is relevant, and the words form a Wordgroup, or whether it is just two words that happen to be next to each other.

Wordgroups are frequently encountered in technical areas, like Medicine (as with “auditory osseointegrated device”), or the law (as with “except as hereinafter provided”). Not all Wordgroups need definitions – for many of them, linking the component words to the exact meaning of the word as it is used in the Wordgroup is enough. Wordgroups are not held in the dictionary, but as a network of objects. Encountering the first word leads to searching for a match. The match may be the first object in a larger Wordgroup (Fig. 4).

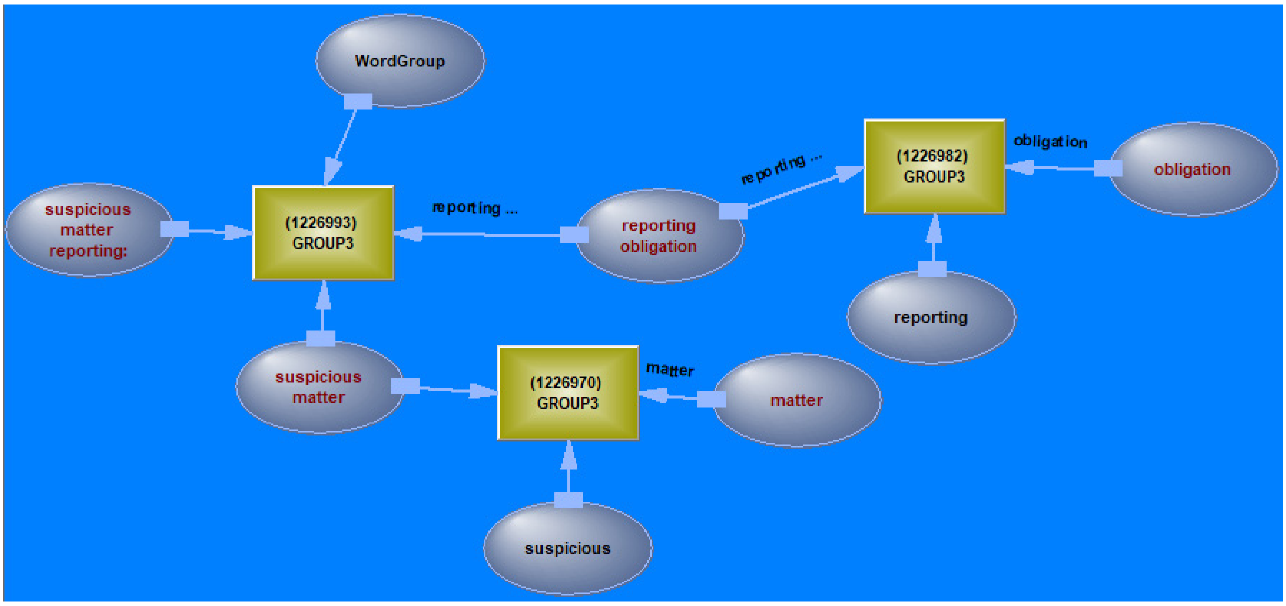

Fig. 4. A sample of the underlying structure used when combining linked words into Wordgroups that can then be operated on as a single unit. This structure represents the Wordgroup "Suspicious matter reporting obligation".

2.3. An Unknown Word

A stable large document can be checked for unknown words before reading. If a less stable document needs to be read immediately, the system may encounter words it doesn’t know. It can look up the word using an API of a human-readable dictionary (HRD), extensively “curate” the entry using instructions in English, then merge it with its other knowledge and continue. Not all words have HRD entries, so the next step is to see if it can be constructed from a root, and prefixes and suffixes. The machine does not use this approach for words with an HRD entry (not all of which are complete), as a constructed word will often acquire new meanings over time ("unbelievable" as an example, where the two meanings are antonymic).

We also need to emphasize that using a dictionary is not the complete answer – it is intended to provide instruction to a person with very powerful innate capabilities for integration of information. Bootstrapping a machine is a different matter entirely. One obvious problem with a dictionary is circularity. For example, the definition of "astonished" is "amazed or surprised", but the definition of "amazed" links right back to "astonished". The dictionary is a human-readable resource – it doesn't need to explain to a person what being amazed or angry or grief-stricken feels like. A person will have experienced them directly, or have seen their effect on others. But we need to explain these things to a machine. Fortunately, the Proof of Concept we have chosen has only one word so far which will need a human explanation – "reckless" – but failure to handle it correctly could result in extremely large fines.

A second example of circularity is the definition of "mend", which links to "fix", which links to "repair", which links right back to "mend". "Put back in working order" is a reasonable alternative to these circular definitions – the result of damage or wear is repaired. An advantage of using English as the control language is that the machine can be tasked to "clean itself up" by finding circularity and other common dictionary faults. A dictionary can also be very shallow – here is the entry for the form of "how" used in the phrase "how to bake a cake" – "in what way or manner; by what means". It needs expansion and clarification – the ingredients don't fit under "means". "Recipe" has a relevant definition, but there is no path through the definitions. There is a path to a clearer definition that can potentially be reached through synonyms, but there is also a great deal of noise that can easily obscure the intended meaning, so we have chosen to avoid loading the HRD synonyms listed with the definitions.

The dictionary may list 20 synonyms for a word – we don't load these unless we have loaded all the meanings for a particular synonym, and can make an exact connection. Linking a specific definition to a word with multiple meanings smears the definition.

For example, suspension is a synonym for abrogation – firstly, it is wrong, as suspension refers to a temporary voiding, while abrogation is permanent, and suspension can be particles in a fluid, or a car's suspension, or a suspension bridge, or prevention of someone exercising their usual role. Unless we can connect the exact definition as a synonym, it is safer not to use them. Later, maybe, when the machine can figure out for itself which of the meanings is meant.

Description: abrogation
Category: Noun
Sense count: 1
Sense 0
Definition: the repeal or abolition of a law, right, or agreement.
Synonym count: 28
…
Synonym 22: abolition
Synonym 23: suspension.

2.4. All the Meanings

It is not just all of the meanings, but all of the parts of speech a word may have that have to be handled (Table 1).

The number of parts of speech is increased because each present participle also has a gerund, each past participle may be an adjective, and each preposition also has a PrepositionVerb, as "He is onto you", or "The police are after him".

Table 1. Breakdown of Our Definition Database with Regard to the Number of Parts of Speech Included for Each Definition. Instances Refers to the Number of Words in Our Set of Entries Which Have the Given Number of Parts of Speech

Instances	Number of parts of speech	Example
3	6	Up – Adjective, Adverb, Noun, Preposition, PrepositionVerb, Verb
35	5	Half
313	4	Sobbing
5277	3	Slashing
12092	2	Lump – Noun, Verb

Definitions are used to help sort out which part of speech is relevant to the instance in question. Determining the appropriate definition is an iterative process during which possibilities are eliminated based on the context in which the instance occurs until only one possibility remains. During this process, connections from Part of Speech through Subcategorization to Definition are tentatively made and when required, broken. If it can't be resolved, it is left until more information is available, or a decision must be made instantly based on the lesser threat.

2.5. Sub Categorization

Subcategorization is used on parts of speech to break them into more manageable groups.

A dictionary does not provide much depth in its handling of verbs. For the most part it classifies them as either intransitive or transitive, with a few notes about multiple objects. Our approach is more thorough, as defining verbs in terms of more discrete categories can aid in the process of whittling down large blocks of definitions to determine an appropriate meaning. Our in-depth approach to categorizing verbs has resulted in close to ninety verb forms (some rare).

1. In "They call the wind Mariah", for example, we classify "call" as a Bitransitive verb, which indicates that it takes two direct objects. Contrast this with "They called the girl Angela to the principal's office", in which case we would classify "call" as a Transitive verb because it takes only a single object. Having the possibility and discounting it allows for a much higher reliability than trying to control the domain so these things do not occur.

To give a more complex example, in the case of a BiTransClausal verb there are two direct objects and a following clause, as is the case with "bet" in the phrase "I bet you any money your horse won't win.

To give two further examples illustrative of the complexity of verb forms:

2. Fred paid John $50 to mow the lawn – ObjectBiTransInfinitive (the first object of "paid" is the subject of the infinitive)

Fred paid John $5,000 to avoid a lawsuit – SubjectBiTransInfinitive (the subject of "paid" is the subject of the infinitive).

Setting aside verbs, even nouns can be somewhat tricky, with a noun having a singular form for one definition, and an aggregate (uncountable) form for another. On top of that, nouns may have any combination of a simple form, an infinitive form, or a clausal form that can support and control a clause, as in "The conjecture that he was drunk was validated." We have thus divided noun usages into the categories: SimpleNoun, InfinitiveNoun, and ClausalNoun, with the ClausalNoun capable of supplying a logical value to the clause it controls.

In the case of adjectives, we have divided up the category into around 25 subcategorizations, including Attributive, Predicative, AttributivePredicative, Clausal, and Infinitive – to name but a few.

We have divided prepositions primarily based on the features that they link, as in NounNoun, ClauseNoun, NounClause, PredicativeNoun, VerbAdjective, and ActionNoun (which may link three objects, as in "he put the money on the table"). Overall, this has resulted in 28 subcategorizations.

2.6. Clumping

Many adjectives occur independently, as with “large” and “black” in the case of “a large black car”. In other cases, multiple adjectives may be linked with one of them serving to modify a structure containing the second, as in “an imaginary flat surface”. In this event, “flat” and “surface” need to be clumped together into an object so that “imaginary” can operate on it (Fig. 5). Note that the node carrying the definition string is also the node representing the structure built from the definition string.

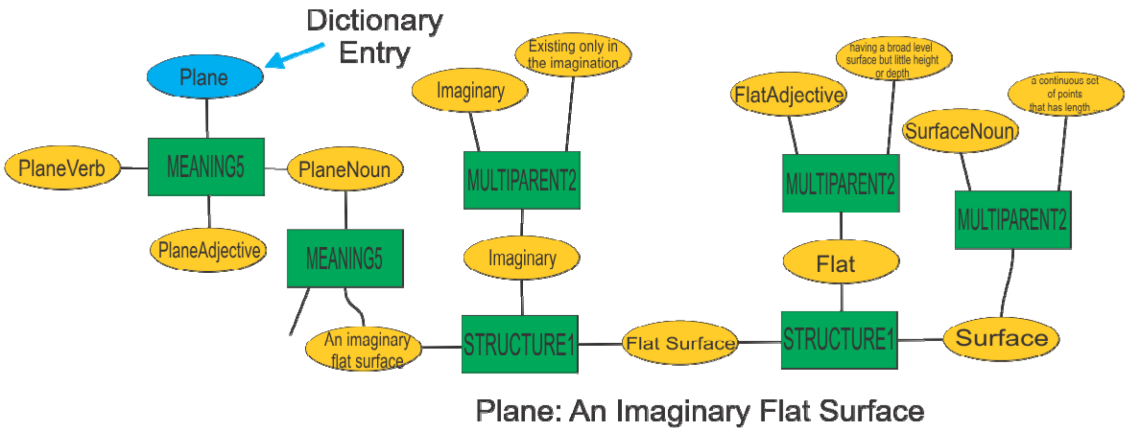

Fig. 5. In the above represented dictionary entry for "Plane", the words "flat" and "surface" have been clumped together into a single structure so that they can be acted upon as a unit.

Prepositions can be clumped with their target subject – “the table in his office” – See Fig. 6.

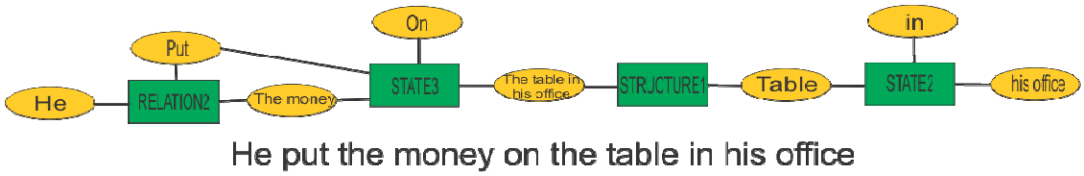

Fig. 6. Illustration of clumping in the case of prepositions.

In the phrase, “He put the money on the table into his bank account”, “the money on the table” becomes an object that can be operated on by another preposition. See Fig. 7.

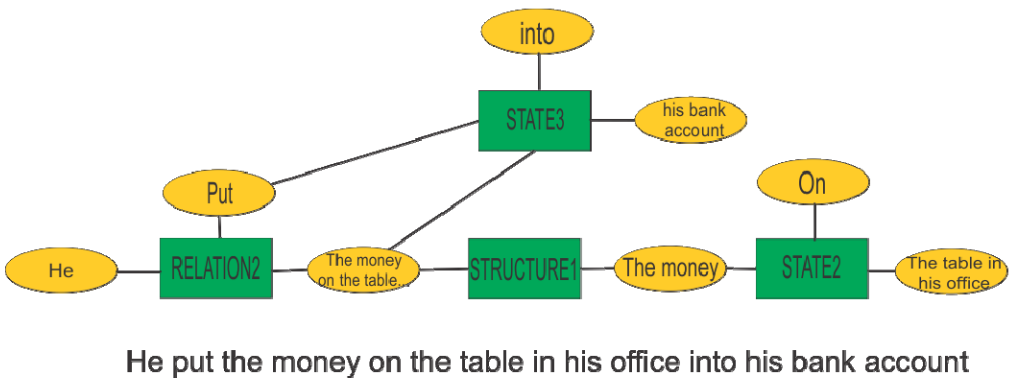

Fig. 7. Illustration of a noun-preposition clump.

2.7. Multiple Definitions

The HRD arranges its definitions in two layers, called Sense and Sub Sense. The Senses are meant to convey independent definitions for a given word, while the Sub Senses are embedded within the Sense and are intended to convey finer or idiosyncratic applications of that Sense. In practice, the relation of Sense to Sub

Sense is not strictly hierarchical, in that something may meet the definition of a given Sub Sense while failing to meet the definition of the overarching Sense. These mismatches of Sense and Sub Sense will often occur when, for example, the Sub Sense represents a metaphorical extension of the idea conveyed by the overarching Sense. These connections may seem intuitive to a human reader, but there is often no way to readily convey this connection to a machine without localized search instructions. This necessitates ignoring the current structure of the dictionary and searching through every possible Sub Sense definition during a meaning search, which will obviously negatively impact overall search speed. Our eventual intent is to allow the machine to rearrange the definitions in an efficient machine interpretable form.

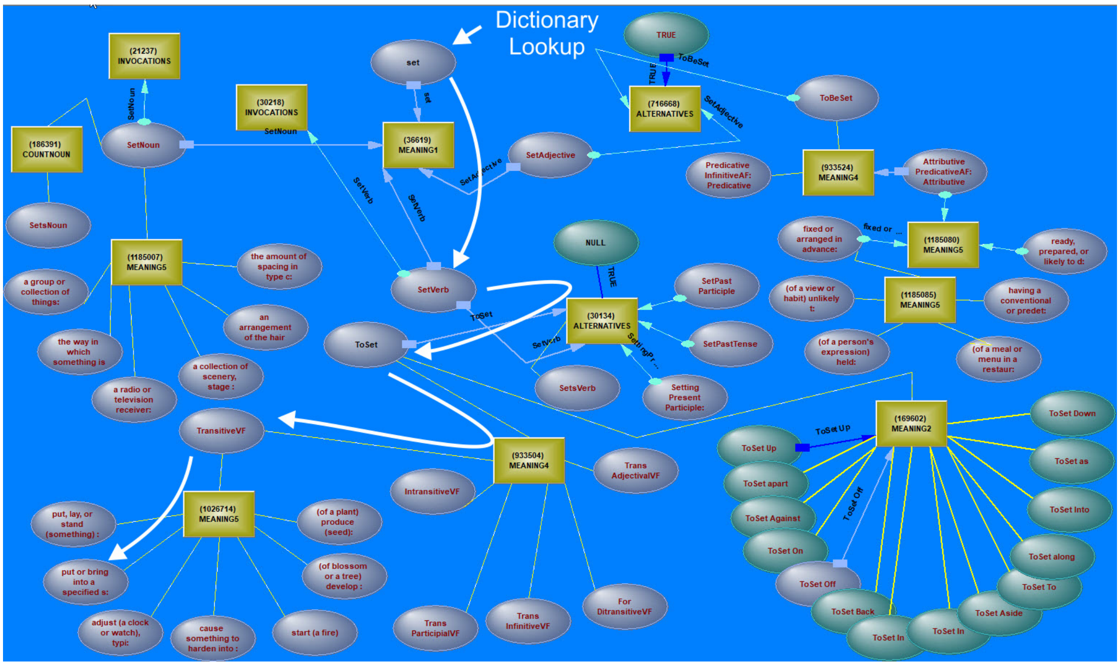

Fig. 8. The word Set and its descendants. The diagram illustrates a small part of a structure with millions of nodes and connections.

Fig. 8 shows part of the Part of Speech, Sub Categorization and definition structure for “set”, together with a collection of collocated verbs for “set” – set in, set up, set off, etc. Each collocated verb has its own definitions (the verb colors the preposition, so in many instances the meaning would not be easily deducible from one of the definitions of the verb and one of the definitions of the preposition). Tagging will bring us to the base word, located in the top center of the figure. Only in the simplest examples is it a sequential process, where we can determine the Part of Speech, then the Subcategorization, then the definition. In most cases it involves an iterative process that takes into account the words around it. When the machine has determined the right combination (it needs a lot of help when starting from nothing), it needs to link the word in the text being parsed to the appropriate base word, a Subcategorization if one exists, and a definition. The child word may inherit properties from the base word and the Subcategorization (Fig. 9), and it inherits the definition (the definition itself has no parents). Some common words (give, open, take) have up to fifteen Sense definitions, and some words have up to fifteen Sub Sense definitions from one Sense definition. In the beginning, the definition is just a node carrying a string holding the curated dictionary definition. Where one word is defined by another single word (which occurs in about 1% of cases), an assignment is used. For other simple cases – for example an ANDing of adjectives or adverbs, the definition node becomes the output of the AND relation. For more involved cases, the definition node becomes an accurate clumping of the words in the definition, as the definition was used to build the structure of the definition.

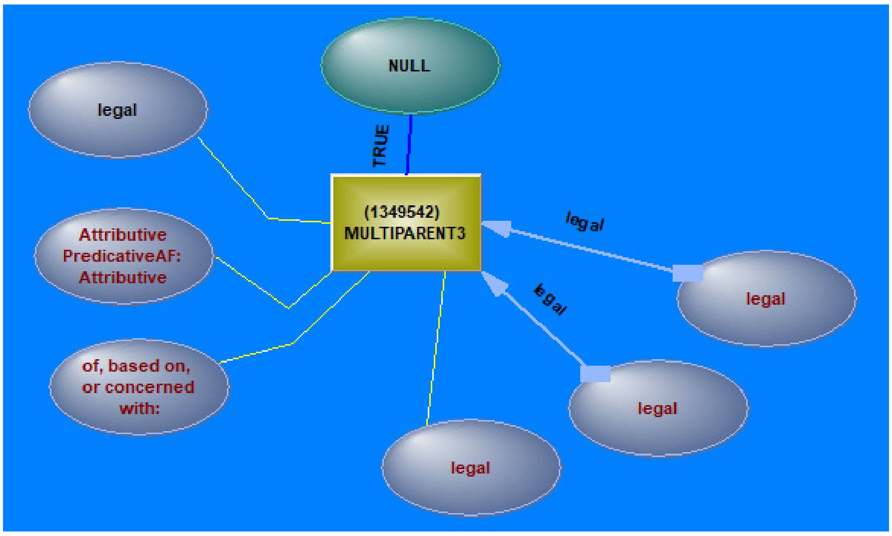

Fig. 9. A MULTIPARENT operator linking its children to multiple points.

Turning words into a working structure aids understanding and troubleshooting (Fig. 10).

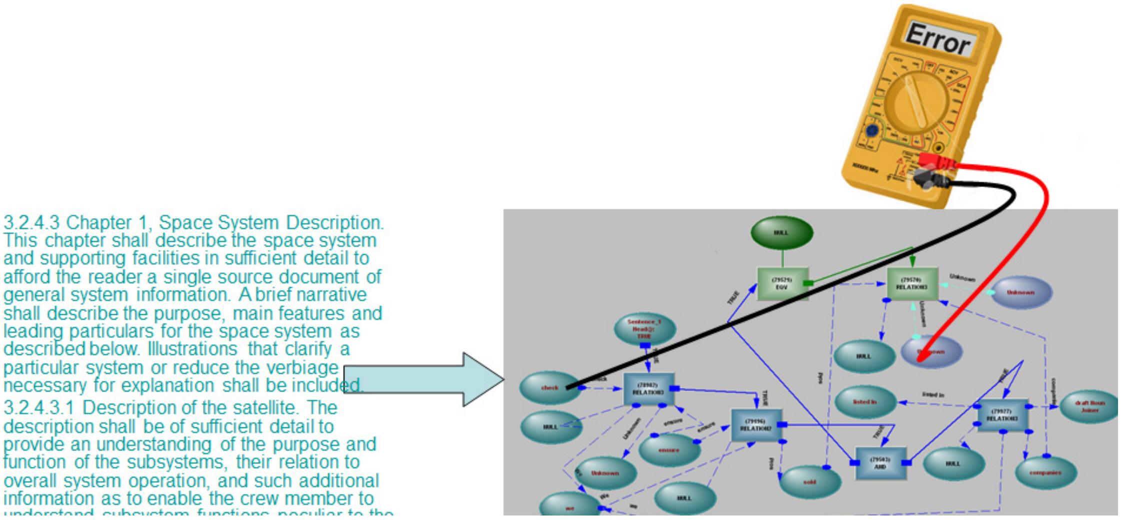

Fig. 10. Troubleshooting a working structure.

2.8. A Reducible Problem?

People have tried to use dictionaries before to handle language from an AGI perspective, but so far the results of these attempts have been negligible. There are several reasons why this would be so.

Compiling a dictionary is about the most reducible problem imaginable. There may be only a few guidelines for those writing the definitions, and thousands of people can contribute to writing those definitions (as is the case with the lexical database Wordnet [27]). Making an automated dictionary reader is a very different problem. Most reasonably difficult problems are reducible – a new aircraft design can be cut into wings, engines, fuselage, control surfaces, avionics, and the groups working on those areas can function largely independently. An automated dictionary reader requires the tight interweaving of many different techniques with repeated iteration, making it much more of a single-mind problem. It is also unlike many problems in that a first-cut solution can still be of some use, and it can be improved over time. Successfully handling text requires many complex things to be simultaneously working smoothly together, otherwise the performance is very poor. A reasonable question would be – "You say it is a hard problem, and you say humans have a limit of four pieces of information in their conscious minds, so why shouldn't we assume that what you have done is a fudge." We have had the luxury of time, forty years all told, to look at different aspects of the problem, to build an operational system for health insurance and observe its shortcomings, and use the unconscious processing ability built up over that long time.

A further question one might ask is, "But why invent all sorts of verb types, and connections for prepositions, when linguists have already done all this?" For two good reasons: trying to describe a system that relies on multiple activities occurring in parallel in an unconscious mind that we don't understand is next to impossible in a static piece of text, and linguists themselves are subject to the very same working memory limitations that serve as an obstacle to complex parallel operations, so their proposed solutions – say frames – will be far from what is required for use in a machine.

3. Anti-Money Laundering as Proof of Concept

Money laundering is the process whereby criminals obscure the source of their illegally-obtained income through stages of placement, layering, and integration [28]. The overall purpose of these actions is to hide the source of the funds through complex financial maneuvering. An estimated $500 billion are laundered every single year [28] and there are cases of financial institutions successfully being sued billions of dollars for their lack of adequate money laundering controls [29]. The overall cost of money laundering in terms of lost revenue easily numbers in the billions, creating a huge demand for the development of AML technologies.

Any involvement of a financial institution throughout the process of money laundering will leave an electronic trail, rendering its detection amenable to direct investigation by AI-technologies. Rapidly combing through bank transactions is easy enough for a machine to accomplish, but this is only part of the solution. Any one bank may only see part of the actual operation, with that part seemingly innocent by itself. A machine attempting to find spurious transactions needs to first determine what to look for and what strategies to pursue. These strategies can be actively decided and programmed by an operator, but this is only an appropriate approach when the problem is a static one and the solution to it is already known, so it won't be used by the money launderers, who learn quickly after witnessing the incarceration of their compatriots.

Combatting money laundering means competing against flexible human agents capable of rapidly changing their methods. The fight against money laundering can best be characterized as an arms race, necessitating constant updating of your strategy to counteract modifications made by your opponent. For this reason, static methods of money laundering detection are of limited usefulness. Despite this flaw, many financial institutions continue to make use of software relying on static, predefined rules [30]. More advanced technologies make use of data mining techniques focusing on anomaly detection [31, 32], social network approaches [33], or the employment of support vector machines (SVM) using probabilistic thresholds [34]. Though these approaches have some merit, their reliance on pre-existing data places a hard limit on their ability to address newly emerging strategies. Furthermore, field testing of these approaches has been limited and their detection-algorithms have been trained using primarily fabricated instances of money laundering [30], casting doubt on their viability in a real-world setting.

A machine capable of recognizing when a strategy must be changed and self-modifying itself in accordance with a new solution will be best equipped to tackle subjects characterized by the sort of variable problem space that typifies the money laundering landscape. The issue is how to recognize that the problem has changed and how to reach a new solution.

Once such a system has been created, resources designed for human use can serve as inputs to help orient the machine in its intended sphere of application. Certain documents designed as aides for human investigation of money laundering offences, such as Australia's Anti-Money Laundering and Counter-Terrorism Financing Act 2006 [35] (henceforth referred to as the AML Act), can serve to set its initial parameters within the sphere of AML. The first step is then for the machine to read and “understand” the AML Act. This translates to mapping out the relational structures present within the Act for future deployment within the problem space. This already has substantial benefits, as a human reader lacking in legal experience can see exactly what is meant by every word and group of words, without getting lost in a forest of bullets and other punctuation.

Our approach relies upon extracting and representing relational structures from relevant blocks of text. The represented relational structures are malleable, activatable and combinable, allowing for incorporation of new information and linkage with other represented structures. The represented structures are highly transparent, and background processes are readily-understandable by non-experts due to direct representation in the English language. Knowledge extension can be readily accomplished by providing access to further sources of information, which will be mapped and incorporated into existing structures. This allows for flexibility when confronted with shifting goals as well as the potential to generate unforeseen solutions to complex problems.

Such an approach can delineate the problem space and identify basic solutions in areas such as anti-money laundering through mapping out the relational structures embedded in legal documents on the subject. The flexibility of the system allows it to evaluate the failings of its own approaches and that of others within the context of the changing AML landscape. The key improvement of our prospective system is its ability to self-modify its strategies in light of new developments in the area of interest such as the advent of cryptocurrency approaches to money laundering.

We hope the following example will serve to illustrate how the objects that words represent can be turned into active machinery.

The AML/CTF Act describes a particular transaction:

Same institution same person electronic funds transfer instruction

(2) For the purposes of this Act, if:

(a) a person (the payer) instructs a person (the ordering institution) to make money controlled by the payer available to the payer by:

(i) being credited to an account held by the payer with the ordering institution; or

(ii) being paid to the payer by the ordering institution; and

(b) the transfer is to be carried out wholly or partly by means of one or more electronic communications; and

(c) the ordering institution is:

(i) an ADI; or

(ii) a bank; or

(iii) a building society; or

(iv) a credit union; or

(v) a person specified in the AML/CTF Rules;

then:

(d) the instruction is a same institution same person electronic funds transfer instruction; and

(e) for the purposes of the application of this Act to making the money available to the payer:

(i) the payer may also be known as the payee; and

(ii) the ordering institution may also be known as the beneficiary institution.

Our intent is that the machine reads the AML Act, and creates a transfer mechanism exactly as the AML Act describes. The mechanism has a name in the format of [Same institution same person electronic funds transfer instruction]. This is treated as a Wordgroup with a definition – the relevant text of the AML Act. The machine already has definitions for a bank account, a credit union, legal niceties ("for the purposes of"), and the collocated verbs used in the text (such as "carried out"). Every word and Wordgroup that is used in this new definition of a more complex object will already have been defined (and the words in their definitions defined).

The machine can then simulate the instruction, by creating a copy of the mechanism and moving money from the account of the payer to the account of the payee, where each object and relation is exactly as the Act describes. This approach eliminates an interpretive step, where a programmer attempts to convert the legal language of the Act to a programming language, leaving much out in the process.

There is risk in every object and relation in the bank's processes – it was once thought that creating a new bank account was a well-controlled process, but that assumption is becoming increasingly untenable. To give a recent example to the contrary, generous flood relief was offered to individuals in affected regions, so someone quickly created 42 fake accounts across 9 banks to take advantage of the situation [36]. A bank account has to be viewed with suspicion – the person, a fake address, or a suspiciously large number of people sharing the address (forty different people applied for flood relief on one house) can be indicators that someone is trying to manipulate the system to their advantage. There are inherent risks in every process on which an enterprising criminal can capitalize.

4. Discussion

4.1. Comparison with Other Technologies

We have waited until after an example of the sort of text that Active Structure is intended to represent has been presented before comparing with other technologies.

4.1.1. Deep Learning – DL

DL is a modification of Machine Learning, where inputs and outputs are stripped away if they make no difference to the output of the device when run against some predetermined block of data which includes all situations the device may encounter. This is a deep, inescapable, and fatal flaw. We are expecting a block of textual data to have answers to problems we have never encountered before? We are choosing to place rigidities into an organization's systems, so that others can exploit them? The results yielded by DL methods thus far have revealed persistent shortcomings, particularly in the interpretation of complex text [37]. When people say "The answer is in the data", this is accurate for DL. If the text is describing something new and different, as legislation or a specification may well be, then DL has no relevance, because the object is not yet in any text from which the data is drawn.

Specifications and legislation describe their own little worlds, with many terms unique to the document defined in open text, in linked documents, or in a glossary, and the text is heavy with bullets and other grammatical structuring to aid human understanding. With its structure represented in a machine, objects that are hundreds of pages away in text may be a link away in the structure (a link that can be activated by the transmission of a state or value). A similar problem exists with Climate Change – the data refers to a state in the past, and a complex of complex models must be tracked instead. A far greater stumbling block is that a word can have many meanings – the word "run" has 82 definitions, the word "set" has 62. To give a few examples, "set" has an entirely different meaning in all of the following instances: a chess set, a movie set, a movie set in Hawaii, the rain set in. Some words even have antonymic meanings – a society has sanctions against murder, while a gang boss sanctions a hit. Or the results were unbelievable – this can mean "extraordinary" or "not believable". For DL methods to be successful, a domain must be found where the words have very few meanings. Such domains, unless they are wholly artificial, are few and far between. Medicine, for example, has cervical vertebrae, cervical cancer, and many other similar collisions. Trying to distil something useful out of a big block of text using an algorithm that has no idea what the words mean is not a promising approach, either for AI or Artificial General Intelligence. The approach guarantees there is no intelligence, no new connection, no reversal of flow. The notion that because an input did not affect the output in the time period, the input can be removed, is dangerous thinking. A fire escape rarely features in the data, but it is proof against catastrophic consequences. The obvious question is “Why do it?”. There is attraction in the notion that a device can be used to avoid thinking about the problem of structure and meaning in language, and this is exactly what DL methods seem to promise. Unfortunately, in the case of hard problems, it doesn’t work at all. Large Language Models (LLM) have the same problem – no matter how many billions of nodes are used, the nuances of language will escape them. Even worse, DL and LLM pay no heed to what was said previously by an interlocutor, so even back fence gossip is beyond them. And for complex specialised text like specifications and legislation, LLM seems to have no place at all.

It seems reasonable to ask, “Dictionaries are built on people reading large slabs of text and accumulating meanings – isn’t this the same thing – using data?” It is different, in that dictionaries list all the meanings of a word, and keep that information up to date as the meanings of words change. Also, it is much easier (but not easy) to choose one meaning from a set of known meanings, instead of trying to find a piece of text that contains the word and other words around it that “match” the words around the particular word being checked. The meanings have been distilled to their essence by human intelligence, and still need extensive curation to suit an intelligent machine.

4.2. Compositionality

Each new development in AI is criticized for lack of compositionality – the words don’t fit together, or act in the way that a native speaker would expect. That is because making them fit together requires much more effort than something which can assemble some images using words as key words.

Compositionality will only be present when all the logical and existential connections are there, all the relations, including the adverbs, are doing what they should.

A small example may indicate the complexity: "He was proud of their success." In this case, “success” controls how proud he should feel, from “not at all proud” to “extremely proud” – the information travels from right to left.

In "He was confident of success”, the information travels from left to right, with “confident” (or “very confident” or “extremely confident”) supplying the logical value for “success”, with the existential value at zero – it hasn’t happened yet. That is, compositionality requires a huge investment in connections and definitions, unrealistic in what is otherwise an amusing toy.

5. Conclusions

It seems we have largely given up the pursuit of AGI, and spend our time on much simpler tasks, or even try to program our way to a data-driven solution using DL, LLM or ML, while deliberately choosing not to analyze the problem.

If we are to handle hard problems using a machine (and now would be a very good time to do so), or use a machine to integrate the contributions of many specialists, we need to make the communication as easy as possible. An obvious way is for the machine to learn English. Machines have had sufficient memory for the last thirty years, so it has only been our lack of understanding of the workings of our unconscious mind that has prevented it. Yes, it is hard to do, because our unconscious mind knows it is too complex for us to understand, so refuses to tell us what it does (it obviously doesn’t do so consciously). Breaking the Four Pieces Limit in a time of great stress seems like a worthy goal.

6. Patents

This work has been proceeding over a long time, beginning in 1982, and everything is strongly interwoven, meaning it would all have to be patented together, or none of it patented. It is far too complicated for a patent.

Author Contributions

Conceptualization, J.B.; methodology, J.B.; software, J.B.; resources, R.S.; writing—original draft preparation, J.B and R.S.; writing—review and editing, J.B. and R.S.; visualization, J.B.; project administration, J.B.; funding acquisition, J.B. All authors have read and agreed to the published version of the manuscript.

Funding

The Australian Government has provided funding based on a percentage of the amount spent by a company on accredited research, for the entirety of its development.

Conflicts of Interest

The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Appendix A: Glossary

Existential Logic: Here, it refers to how existence is described in English – "there", the verb auxiliaries "can", "able", etc., the suffix "-able", and how an existential state can be combined with a propositional or temporal state.

Modus Tollens: A logical argument which takes the form of "If P, then Q. Not Q. Therefore, not P." In a network, not P would be radiated.

Self-Extensible: Describes a structure which has resources available to it and can create new structure depending on circumstances in the network. In particular, a structure which can make new connections, activate them, evaluate the result, and either revert to a previous state or continue in the new state.

Undirected Network: Here meaning a network in which information can flow in any direction, controlled by decisions made by the operators in the network rather than some general constraint.

References

[1] Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, *30*, 191–238. doi: 10.1016/0010-0277(88)90020-0.

[2] Pinker, S. (1994). *The Language Instinct* (pp. 228-229). HarperCollins Publishers: New York, USA.

[3] Cacciari, C., & Tabossi, P. (1993). *Idioms: Processing, Structure, and Interpretation*. Psychology Press: New York, USA.

[4] Pritchett, B. L. (1992). *Grammatical Competence and Parsing Performance*. Chicago: University of Chicago Press: Chicago, USA.

[5] Jackendoff, R. (1994). *Patterns in the Mind* (pp. 66-82). Basic Books: New York, USA.

[6] Falkum, I., & Vicente, A. (2015). Polysemy: Current perspectives and approaches. *Lingua, 157*, 1-16.

[7] Bergenholtz, H., & Gouws, R. (2017). Polyseme selection, lemma selection and article selection. *Lexicos, 27*, 107-131, doi: 10.5788/27-1-1396.

[8] Tanenhaus, M. K., & Trueswell, J. C. (1995). Sentence comprehension. In Miller, J. L., Eimas, P. D. (Eds.), *Speech, Language and Communication* (pp. 217-262). 1st ed., New York: Academic Press

[9] Seidenberg, M. S., Tanenhaus, M. K., Leiman, J. M., & Bienkowski, M. (1982). Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing. *Cognitive*

Psychol, 14, 489-537. doi: 10.1016/0010-0285(82)90017-2.

[10] Singh, C. P., Devi, H. M., & Sing, N. A. (2021). A review on word sense disambiguation emphasizing the data resources on Wordnet and Corpus. *ITII*, *9*, 996-1016, doi: 10.17762/itii.v9i2.444.

[11] Lee, C., & Federmeier, K. D. (2012). Ambiguity's aftermath: How age differences in resolving lexical ambiguity affect subsequent comprehension. *Neuropsychologia, 50*, 869-879, doi: 10.1016/j.neuropsychologia.2012.01.027.

[12] Lee, C., & Federmeier, K. D. (2011). Differential age effects on lexical ambiguity resolution mechanisms. *Psychophysiology, 48*, 960-972, doi: 10.1111/j.1469-8986.2010.01158.x.

[13] Gadsby, N., Arnott, W. L., & Copland, D. A. (2008). An investigation of working memory influences on lexical ambiguity resolution. *Neuropsychology*, *22*, 209-216, doi:10.1037/0894-4105.22.2.209.

[14] Cowan, N. (2010). The magical mystery four: How is working memory capacity limited, and why? *Curr Dir Psychol Sci, 19*, 51-57, doi: 10.1177/0963721409359277.

[15] Altmann, G. (2001). The language machine: Psycholinguistics in review. *Br J Psychol, 92*, 129-170, doi:10.1348/000712601162130

[16] Mitchell, T., Cohen, W., Hruschka, E., Talukdar, P., Yang, B., Betteridge, J., & Welling, J. (2015). Never-ending learning. *Proceedings of the National Conference on Artificial Intelligence* (pp. 2302-2310), doi: 10.1145/3191513

[17] Dupoux, E. (2018). Cognitive science in the era of artificial intelligence: A roadmap for reverse-engineering the infant language-learner. *Cognition, 173*, 43-59, doi: 10.1016/j.cognition.2017.11.008.

[18] Carlson, A., Betteridge, J., Kisiel, B., Settles, B., Hruschka Jr., E. R., & Mitchell, T. M. (2010). Toward an architecture for never-ending language learning. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence* (pp. 1306-1313).

[19] Mu, R., & Zeng, X. (2019). A review of deep learning research. *KSII Trans Internet Inf Syst, 13*, 1738-1764, doi:10.3837/tiis.2019.04.001

[20] Otter, D. W., Medina, J. R., & Kalita, J. K. (2019). A survey of the usages of deep learning for natural language processing. *IEEE Trans Neural Netw Learn Syst, 32*, 604-624, doi:10.1109/TNNLS.2020.2979670

[21] Marcus, G. F. (2018). Deep learning: A critical appraisal. *arXiv*1801.00631, doi:10.48550/arXiv.1801.00631

[22] Goertzal, B. (2014). Artificial general intelligence: Concept, state of the art, and future prospects. *JAGI, 5*, 1-48, doi:10.2478/jagi-2014-0001

[23] Mikolov, T., Joulin, A., & Baroni, M. (2016). A roadmap towards machine intelligence. In *International Conference on Intelligent Text Processing and Computational Linguistics,* 29-61, doi:10.1007/978-3-319-75477-2_2.

[24] Active Semantics. *Lane Following*. From https://activesemantics.com/wp-content/uploads/2022/10/Lane-Following1.pdf

[25] Marcus, G. F. (2003). *The Algebraic Mind: Integrating Connectionism and Cognitive Science*. MIT Press: Cambridge, USA.

[26] LeCun, Y., & Browning, J. (2022). What AI can tell us about intelligence. In *NOEMA Magazine.* From https://www.noemamag.com/what-ai-can-tell-us-about-intelligence

[27] Miller, G. A. (1995). Wordnet: A lexical database for English. *Commun ACM, 38*, 39-41, doi:10.1145/219717.219748.

[28] Salehi, A., Ghazanfari, M., & Fathian, M. (2017). Data mining techniques for anti-money laundering. *Int J Appl Eng, 12*, 10084-10094.

[29] Huang, J. Y. (2015). Effectiveness of US anti-money laundering regulations and HSBC case study. *J Money Laund Control*, *18*, 525-532, doi: 10.1108/JMLC-05-2015-0018.

[30] Chen, Z., Van Khoa, L. D., Teoh, E. N., Nazir, A., Karuppiah, E. K., & Lam, K. S. (2018). Machine learning techniques for anti-money laundering (AML) solutions in suspicious transaction detection: A review. *Knowl Inf Syst*, *57*, 245-285, doi: 10.1007/s10115-017-1144-z.

[31] Ahmed, M., Mahmood, A., & Islam, M. (2016). A survey of anomaly detection techniques in financial domain. *Future Gener Comput Syst, 55*, 278-288, doi:10.1016/j.future.2015.01.001.

[32] Le Khac, N. A., & Kechadi, M. T. (2010). Application of data mining for anti-money laundering detection: A case study. In *Proceedings of the IEEE International Conference on Data Mining* (pp. 577-584), doi:10.1109/ICDMW.2010.66.

[33] Fronzetti C. A., & Redmondi, E. (2017). Using social network analysis to prevent money laundering. *Expert Syst Appl*, *67*, 49-58, doi:10.1016/j.eswa.2016.09.029.

[34] Kingdon, J. (2004). AI fights money laundering. *Intell Syst, 19*, 87-89, doi:10.1109/MIS.2004.1.

[35] AUSTRAC. Anti-Money Laundering and Counter-Terrorism Financing Act 2006. From https://www.austrac.gov.au/business/legislation/amlctf-act

[36] Daily Mail Australia. From https://www.dailymail.co.uk/news/article-11109469/Bill-Shorten-lashes-NSW-Queensland-flood-claim-frauds-scumbags.html

[37] Li, H. (2018). Deep learning for natural language processing: Advantages and challenges. *Natl Sci Rev, 5*, 24-26, doi:10.1093/nsr/nwx110.