Research Article ISSN 3069-8316

Journal of Advances in Artificial Intelligence and Machine Learning

LLMs, Neurosymbolics, Semantic AI. and Reality

Jim Brander*

Interactive Engineering, Sydney, NSW, Australia.

*Correspondence:
Jim Brander, Interactive Engineering, Sydney, NSW, Australia, Email: jim.brander@activestructure.com.au

Received: 16 Dec 2025; **Accepted:** 29 Jan 2026; **Published:** 08 Feb 2026

Citation: Jim Brander. LLMs, Neurosymbolics, Semantic AI. and Reality. J Adv Artif Intell Mach Learn. 2026; 2(1): 1-4.

ABSTRACT

A comparison of how LLMs, Neurosymbolics and Semantic AI handle a piece of legislation running to 400 pages. Its structure is typical of legislation, specifications, legal agreements, contracts, where a misunderstood word may cost millions to billions, or someone's life.

Keywords

Large Language Models (LLMs), Neurosymbolic AI, Semantic AI, Hybrid AI systems.

Introduction

The document we will be using for test purposes is Australia's Anti-Money Laundering and Counter-Terrorism Financing. It is 400 pages long, and is written in a mixture of Legal and Business English (LBEnglish), with an index structure. It has:

270 internal links (following the link also means changing the logical environment as one does so)

200 external links to other nominated legislation – Crimes Act, Corporations Act, Criminal Code

2 Meta-instructions - ***security*** has the meaning given by section 92 of the *Corporations Act 2001* (for this purpose, disregard subsections 92(2A), (3) and (4) of that Act).

A concerning phrase:
10 links to **any other law** of the Commonwealth, State or Territory
A good reason why there should be mechanisation within the library.

We will use the noun "bar" as an example of a word with many senses – it has a diversity between physical, abstract and figurative objects, and as a collective noun ("bar" can also be a verb or a preposition).:

Iron bar – a physical object that can be used to block access
A bar of music/light/colour/soap/chocolate – "hum a few bars"
A colour bar
Wine bar – an establishment providing alcoholic drinks
A bar – an abstract prohibition
Raise the bar – a figurative allusion to the horizontal bar on a high jump frame – raising it makes it harder to jump over.
A bar – a unit of atmospheric pressure
The Bar – a collective noun for lawyers

Two test sentences – a figurative allusion and an abstract prohibition:

TSMC raised the bar on semiconductor track widths at the conference with its use of Far Ultraviolet.
Fred raised the bar on forever chemicals in the drinking water at the meeting.

To tell these apart, the reader would need to know that "forever chemicals" is a synonym of PFAs, and PFAs above a certain concentration are banned from the water supply (in other words, a human has to know a lot to function in society – the hope that machines not knowing much can solve problems beyond human ability seems fanciful – they are going to have to know a great deal).

LLMs

LLMs were invented for Search Engines. It was then found that simple reasoning based on word association could be handled. Abramski et al. [3] describes word association. The method has several weak spots:

A. words can have several parts of speech;
He heard a car approaching, and turned on the property's security lights. The car turned on a dime and disappeared.

B. words can have multiple meanings (see "bar") - the LLM can't be sure the senses are the same, because it doesn't know what the senses are;

C. complexity.
As complexity grows, the context moves further away from the object (that is, the context needed may be tens or thousands of pages away from the object, and links (such as "A defendant bears an evidential burden in relation to the matter in subsection [2]" are necessary. Following the link means validating the logical environments at both ends of the link are compatible – a far cry from simple word association.

Shojaee P, et al., [3] give a thoughtful account of why LLMs are not good at reasoning, but not knowing what words mean seems a simpler explanation.

Yes, an LLM can save a great deal of time and effort, but it is unreliable in a complex environment.

It may be possible to cobble together something that masks its weaknesses – a Python graph when it must know what words mean, something for meta-instructions, but a cobbled together system has a high risk of failure

Neurosymbolics

Neurosymbolics is an amalgam of Artificial Neural Nets and Symbolic Logic. Symbolic Logic can be used to derive much of mathematics. Gary Marcus gives a strong defence of NeuroSymbolics in [4].

Its Achilles Heel is that it operates on symbols, and it requires a black and white world where everything is unequivocally true or false.

"This means that it is true in all possible worlds and under all interpretations of its non-logical terms, like the claim "either it is raining, or it is not"".

This sounds great – the purity of mathematics in AI. Let's take the claim "It is raining or it is not". It is called **virga**, when rain falls from a cloud, but evaporates before reaching the ground. "Is it raining?" Hard to say. The restrictions are so great that Symbolic Logic may be useful for mathematics, but nowhere else, including science, certainly not as a waystation to AGI.

People cannot be turned into symbols – they have an inner life – they have intelligence, they have moods, emotional drivers, concepts of honour, decency (or not). People can be bad actors, working clandestinely or openly to destroy what other people are working towards. Small children in particular are deemed "unpredictable" near traffic, as they will show agency and run out onto a road, as they haven't learned to fear being run over. Other organisms, from viruses to higher animals, have means of protecting themselves, so trying to make them symbols does not work well. Some operations – a caterpillar turning into a butterfly or a chess pawn turning into a queen, are just not mathematical and shouldn't be anywhere near Symbolic Logic.

Natural phenomena – a line from a meteorologist – "bushfires create their own weather". Bushfires can be predicted to some extent, but not easily represented by a symbol, given their ability to break out of containment. Neurosymbolics is not good at handling figurative allusions, metaphor or idiom – it is happy in the simple world of $X = Y + Z$. A system which can't see further than the mathematics is likely to be a nonstarter in almost every AGI problem. A recent example – a bridge was built on a high mountain in China. A bridge pylon was founded on a slab of rock lying at an angle to the horizontal. The rock moved, and the bridge tumbled down into a gorge. Successfully solving difficult problems requires an understanding of the environment (physical and human) in which the solution will be placed – a one-trick pony will rarely do.

There is a naivety in Computer Science that says problems can be solved with symbols having exactly one meaning – programming. Simple problems – yes. Complex problems - no. Complex problems need a complex language to describe them – a natural language that is continually kept up to date by having to describe advances in every field, from quantum entanglement to genetics.

Semantic AI

Wikipedia [5] provides an excellent introduction to Natural Language Processing across a very wide field. We have a much narrower focus – the turning of a large piece of text, such as a piece of legislation or a specification, into an active structure. Humans have a severe limit on their Conscious Mind – a limit of no more than four things in play at once.

The problem shows up in:
Legislation – suicides, billions in reparations (Robodebt – a 5-second analysis would have stopped it dead)
Specifications – billions wasted (Constellation), hundreds of billions wasted (F-35)
Process Manuals – 2,000 pages (also Constellation)
Physical Modelling – Climate Change – no figure provided, but surely in the trillions
Collaboration – a dozen subject matter experts and a lawyer – a recipe for failure
Economic Modelling – there is a push to use Complex Systems Engineers, as economists are not good at
economics
Bad actors – Volkswagen and exhaust emissions, Boeing and FAA regulations
Money Launderers – nimble and quick-witted adversaries against slow-moving banks and corporations
Lots of others

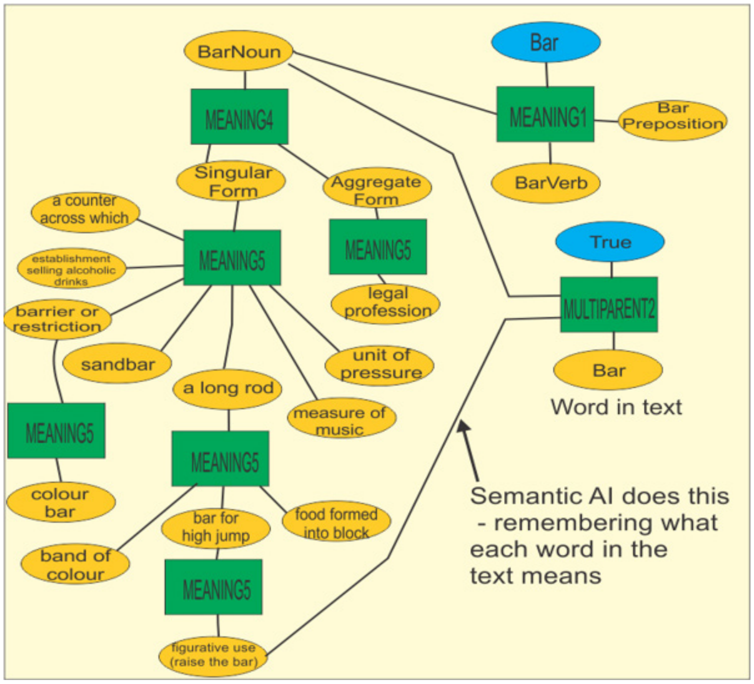

Figure 1: Senses of the noun Bar.

The goal of Semantic AI is to be at least as competent as a well-educated human in understanding English – particularly LB English – with its words having many meanings, its figurative allusions ("raise the bar", "fishing for compliments"), its metaphors, its idiom, and far more competent at handling large, complex texts (humans run out of puff at 50 pages, legislation often runs to 1,000 pages, some specifications run to 80,000 pages – humans tackle them as part of their job, but often make an awfully expensive mess – not their fault – they need tools). When the Semantic AI system encounters a word with multiple meanings ("bar", say), it attempts to use local context or distant linked context to resolve the meaning. It will, if necessary, reduce the number of meaning possibilities and mark for return later (that is, most of the effort is to write the meaning of the word on its back – see Figure 1). It has to find a coherent solution across hundreds or thousands of things – something we can't do. No time is spent in trying to translate into a much more limited language, where meanings would be lost or misconstrued. What's the cost? The machine has to do all the work that your Unconscious Mind does to prepare the text Where it goes further is that objects acquire attributes (a car has a body colour, a motor type and a traction battery – if fitted), an operator has a means of starting, providing fuel to, and stopping a motor. A person can be calm, enraged or depressed (or lots of intermediate states, so an avatar can approximate reality). The effort is worthwhile – if the text is the specification for an aircraft, add a few formulae or graphs, and the aircraft is ready to fly (in an abstract sense). Start the engine, let go the brakes, the plane lifts off, stow the undercarriage and climb. The network structure containing objects and operators is undirected and self-extensible – the fuel flows to the engine, the force flows to the airframe. Change the simulation of a ten-hour mission to last 5 minutes. The machine doesn't have to reason about how something works, it can watch its performance, catch any shortcomings. If the design isn't complete, it (or you) can try to find something that works, or works better while the design is still "on paper". The trade-offs become agonisingly real. Component commonality is something that could be kept more under control at the paper stage, because more conflicts would be resolved.

The same approach works for legislation, although with Anti-Money Laundering, you are probably better off making the machine find a way around the obstacles in the legislation – it can simulate thinking like a crook too. Given the recent past, this will be an important part of AGI. It is also why who uses such a tool becomes very important.

Why Use LBEnglish as the AI Language?

A better question might be: Why not use LBEnglish? If you only deal with short, simple documents or programming, you can skip the extra work that LBEnglish as an AI Language poses. If the document is to be used by many people over a long period, or a mistake could be costly (millions) or embarrassing, then the effort is worthwhile.

What about human involvement?

Human involvement starts with the handler – preferably someone trained in Complex Systems Engineering with experience in dealing with upper management.

There will be layers of management that need to be convinced – the machine has to get some runs on the board. That is why a considerable effort has been put into understanding the drivers of a human – pride, status, self-image, and the emotional states – the shades of anger, rage, jealousy, depression (about 300 in all). Just providing an elegant solution to a hard problem will be nowhere near enough – the machine is going to have to sell it to a sceptical, if not occasionally hostile, audience.

References

1. Jim Brander, Ryan Sigmundson. An Active Structure-Based Artificial General Intelligence. Journal of Advances in Artificial Intelligence. 2023; 1: 1-19.
2. Abramski K, Improta R, Rossetti G, et al. The LLM World of Words English free association norms generated by large language models. Sci Data. 2025; 12: 803.
3. Shojaee P, Mirzadeh I, Alizadeh K, et al. The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity. 2025. arXiv:2501.12948.
4. Marcus G. NeuroSymbolics. Wikipedia.
5. https://en.wikipedia.org/wiki/Natural_language_processing

Appendix

Some numbers:

Vocabulary	45,000	
Definitions:	100,000	
Wordgroups	10,500	Verb modifies preposition – “thin out”
		Figurative – “raise the bar”
		Idiomatic – “kick the bucket”
		Proverb – “a stitch in time”
Nouns	16,000	
Verb forms	110	Example: BiTransClausal – “I bet you any money your horse won’t win”.
Verbs	8,500	(which become operators in active text)

Synonyms There are 80,000 synonym links – they come from a particular sense of the word, but the dictionary does not control the sense of the word where they go. Working on it.