

Tax, LLMs, RAGs and some soft logic

Bernie Walrut*

Abstract

This article explores the usefulness of prompt and Retrieval Augmented Generation (RAG) enhanced large language models (LLMs), a form of artificial intelligence, in tax practice. It analyses the advice provided by selected LLMs in respect of three Australian tax problems, with and without Prolog-like rules in the prompt and in localised RAG files. The analysis highlights that without such enhancement the output from LLMs regarding the application of detailed tax rules is of limited use. However, with Prolog-like enhanced prompts, some models provided useful draft advice. The RAG file approach was less useful. The article then discusses how the tax professional might utilise such tools.

Keywords: large language models, utility, tax practice

* LLB (University of Adelaide); GD Aq (Deakin University); M Aq (Deakin University); MIT (University of South Australia); MTax (University of New South Wales); LLM (University of Queensland); PhD (University of British Columbia). Barrister, Murray Chambers, Adelaide. Email: bwalrut@bpwc.biz.

1. INTRODUCTION

The field of artificial intelligence (AI) has recently experienced a remarkable resurgence.¹ This resurgence has been largely driven by advancements in generative AI technologies, particularly the emergence of large language models (LLMs) and large multimodal models. This is progressing into the development of large action models,² models that go beyond language-based assistance and can trigger events. There is also a developing body of small language models that can equal or surpass LLMs in some fields.³

Whilst there is extensive publicity about the possible use of LLMs in many fields of endeavour including the law, there is less general discussion about the use of LLMs in day-to-day tax practice. Much of the discussion that is available appears to fall into three main areas: the use of AI by taxation authorities to improve the efficiency and effectiveness of their processes, more general discussion about its impact in policy and the imposition of tax, and a historical discussion about the use of expert systems. In day-to-day tax practice LLMs are being used to assist practitioners in, inter alia, preparation of documents, all forms of written communications, correcting grammar, summarising documents, assisting in research and case analysis,⁴ creating chronologies, court discovery processes and comparing documents.⁵

This article seeks to explore the use of LLMs in tax practice in preparing advice on tax outcomes and/or tax consequences of a given fact situation.⁶ However, the focus of the article is not simply on exploring the use of LLMs and apps that allow the tax practitioner to provide such advice; it also includes a consideration of the use of prompt libraries and LLMs suitably enhanced by the use of local Retrieval Augmented Generation (RAG) files.

In the course of that exploration, the article compares the results of posing three types of queries to a number of LLMs. The first was a query with the basic facts and outcome sought. The second was a query with a prompt enhanced by the relevant rules and where

¹ The modelling and analysis underlying this article were undertaken in December 2024 and January 2025 using the LLM models and applications then available. The article was subsequently prepared, drawing on that analysis, for presentation at the 16th International Atax Tax Administration Conference held in Sydney on 8-9 April 2025. Following feedback from the conference, the article was revised and submitted for publication in May 2025. Since the original analysis was undertaken, many new models have emerged, and the applications used have undergone significant improvement. The law publishers' products, as well as commercially available legal tools, have also improved significantly. This article has been revised to give effect to some of the suggestions of an anonymous reviewer, for which the author is most grateful. All errors and oversights remain the responsibility of the author.

² Lu Wang, Fangkai Yang, Chaoyun Zhang, Junting Lu, Jiaxu Qian, Shilin He, Pu Zhao, Bo Qiao, Ray Huang, Si Qin, Qisheng Su, Jiayi Ye, Yudi Zhang, Jian-Guang Lou, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang and Qi Zhang, 'Large Action Models: From Inception to Implementation' (2025) *arXiv:2412.10047*.

³ Xinyu Guan, Li Lyna Zhang, Yifei Liu, Ning Shang, Youran Sun, Yi Zhu, Fan Yang and Mao Yang, 'rStar-Math: Small LLMs Can Master Math Reasoning with Self-Evolved Deep Thinking' (2025) *arXiv:2501.04519*.

⁴ In Australia there are no commercial tax-specific research LLMs, such as the American and Canadian LLM, Blue J. See Blue J, 'Tax Research You Can Rely On' <<https://www.bluej.com>>.

⁵ Mathew Dahl, Varun Magesh, Mirac Suzgun and Daniel E Ho, 'Hallucinating Law: Legal Mistakes with Large Language Models Are Pervasive', *Stanford University Human-Centered Artificial Intelligence* (11 January 2024) <<https://hai.stanford.edu/news/hallucinating-law-legal-mistakes-large-language-models-are-pervasive>> (accessed 18 January 2025).

⁶ Much the same process may be utilised by tax administrators in tax compliance programs.

required tax rates. The third query embodied the rules (and rates where required) in a RAG file, which was used to enhance the model. In the case of the enhanced prompt and the RAG files, the relevant tax rules were expressed in a form recognisable by a non-algorithmic computer language that has often been used in expert systems, Prolog. In this article such rules and facts are described as Prolog-like rules.⁷

As will be explored briefly in this article, Prolog has been important in the development of expert systems. Also, as will be demonstrated, the enhancement of the LLMs was required to navigate the intricate regulatory landscapes and complex fact situations often associated with providing and preparing advice on tax outcomes and/or tax consequences. This field demands not only precise information retrieval but also a deep contextual understanding of complex tax rules and often complex fact situations.

Prior to discussing the outputs of these LLMs, this article provides a brief history of AI, focusing on early expert systems. Those systems were heavily influenced by methodologies such as first-order predicate logic, non-algorithmic languages and non-coding methods including decision trees and graphs.

2. SOME HISTORY AND CONCEPTS

2.1 AI, machines and languages

The possibility of developing machines that replicate human (thought-like) processes has long been dreamed about and discussed.⁸ With the development of modern machines has come the development of high-level computer languages that provide the programmer with a language more like human language and free the programmer from the need to undertake programming with intricate steps and memory management.⁹

Many AI commentators describe in some detail Alan Turing's work from the 1940s, how the 1956 Dartmouth Summer Workshop at Stanford University 'marked the beginning of the artificial intelligence enterprise', and how that workshop name, 'artificial intelligence'¹⁰ was adopted.

George Luger provides an overview of the areas in which AI can be applied as: game playing, automated reasoning and theorem proving; expert systems, natural language understanding and semantic modelling; modelling human performance; planning and robotics; and machine learning.¹¹

⁷ The rules and facts in that form were run in SWI Prolog using Swish at <<https://swish.swi-prolog.org>> and the runs produced the expected answers. However in this article the emphasis is not on the programming but expressing the rules and facts in a soft logic form. The reason for using Prolog-like rules is that they express the relationships more simply and in a manner long used in expert systems, as will be described.

⁸ George F Luger, *Knowing Our World: An Artificial Intelligence Perspective* (Springer, 2021) ('*Knowing Our World*'). This discussion does not consider the early discussions about natural language modelling and processing.

⁹ Ibid 14. These languages include C, Java and Python.

¹⁰ See ibid 58-60 on the derivation of that expression.

¹¹ George F Luger, *Artificial Intelligence: Structures and Strategies for Complex Problem Solving* (Addison Wesley, 4th ed, 2002) ch 2 ('*Artificial Intelligence*').

Notwithstanding periods of intense interest, there have also been long periods when AI research and development has dwindled, sometimes described by the AI community as AI winters.¹²

The texts in the area commonly break the elements of AI and machine learning down into: problem solving (commonly using various search techniques); knowledge representation, reasoning and planning; the use of model-based reasoning and solving under uncertainty and over time; learning, including from examples, using probabilistic models and reinforcing the learning; and communicating, perceiving and acting including the use of natural language processing, communication and recognition models.¹³ One example of problem solving using model-based learning is the use of expert systems.

In the broader commercial world (including in legal, accounting and tax practice), machine learning is widely used in software to enable an improved experience for the user, with some software learning about the user's behaviour after its first use. After the software has been in use for a period of time it begins to predict what the user wants to do.¹⁴ In legal practice the use of AI appears to have focused, so far, on the areas of document assembly, due diligence, contract analysis, automated discovery, natural language searching and attempting to predict the outcome of litigation.¹⁵

2.2 Expert systems

2.2.1 Introduction

At its simplest, an expert system is '[a] computer program designed to model the problem-solving ability of a human expert'.¹⁶ The General Problem Solver by Newell, Shaw and Simon developed between 1955 and 1959¹⁷ is usually described as the first attempt to separate problem-solving methods¹⁸ from knowledge of the problem.¹⁹ Expert systems have historically had two elements: knowledge from a human expert²⁰ and code that puts this knowledge in a form that a computer may apply to solve similar problems.²¹ Some add the interface or other aspects of the computing system as a third

¹² Luger, *Knowing Our World*, above n 8, 223; David C Brock, 'Learning from Artificial Intelligence's Previous Awakenings: The History of Expert Systems' (2018) 39(3) *AI Magazine* 3, 4.

¹³ Luger, *Artificial Intelligence*, above n 11.

¹⁴ Jason Bell, *Machine Learning: Hands-On for Developers and Technical Professionals* (Wiley, 2014) 4.

¹⁵ See AI Legal Assistant, 'The Gold Standard in Legal Productivity' <<https://legalassistant.au>> (accessed 24 December 2024) for a longer list.

¹⁶ John Durkin, *Expert Systems: Design and Development* (Prentice-Hall, 1994) 7. Some prefer to call them knowledge-based systems. See Brock, above n 12, 14; Xianzhe Yang and Changsheng Zhu, 'Industrial Expert Systems Review: A Comprehensive Analysis of Typical Applications' (2024) 12 *IEEE Access* 88558.

¹⁷ A Newell, JC Shaw and HA Simon, *Report on a General Problem-Solving Program* (Rand Corporation, 1958, rev 1959) available at: <https://archive.org/details/bitsavers_randiplP15roblemSolvingProgramFeb59_1757450/page/n3/mode/2up>.

¹⁸ A logic-based or heuristic search. See Brock, above n 12, 5.

¹⁹ Luger, *Artificial Intelligence*, above n 11, 20; Durkin, above n 16, 7; Haocheng Tan, 'A Brief History and Technical Review of the Expert System Research' (2017) 242 *IOP Conference Series: Materials Science and Engineering* 012111, 1; Brock, above n 12, 5. As Brock (at 8) describes, this became known as knowledge engineering.

²⁰ Referred to as domain knowledge or the knowledge base, which involves storing the facts and the rules.

²¹ Luger, *Artificial Intelligence*, above n 11, 20; Durkin, above n 16, 7; Yang and Zhu, above n 16, 88559.

element.²² Most expert systems involve a particular representation of modelled knowledge, known as production rules, the ‘if then statements’,²³ or the ‘conditional+result’ structure.²⁴ The expert knowledge is a combination of a theoretical understanding of the problem and a collection of heuristic problem-solving rules that experience has shown to be effective in the domain of knowledge.²⁵

A very recent review of the use of expert systems in the industrial field highlights that recent innovations in the use of convolutional neural networks and support vector machines may be incorporated into expert systems to ‘offer more intelligent and efficient decision support’.²⁶ It emphasised that some of the more modern systems use learnt outcomes or other sources of information to assist with the reasoning²⁷ in addition to or beyond that provided by the expert.

Most expert systems to date have been written for specialised expert-level domains and clearly defined problem-solving strategies. They have been applied in many areas including medicine, education, business, design and science. Wagner, in a 2017 article using qualitative analysis of a large body of expert systems cases in the literature, described expert system research as one of the ‘longest-running, and most successful areas of ongoing research within the AI field’ since the 1980s.²⁸ Wagner described the use of expert systems by industry,²⁹ showing that the use of expert systems by legal services in the population studied was very small (constituting 0.6 per cent of the total use across the industries described) and in accounting services it was much more significant (constituting 10.9 per cent).

One of the advantages of expert systems is that they produce more consistent results than humans.³⁰ They have also been used to assist human experts, by assisting with routine tasks, assisting in managing complexities and assisting in addressing issues that may not be readily recalled.³¹ They often provide precise solutions.³²

Yang and Zhu recently described some of the limitations of modern expert systems, such as being dependent on expert knowledge, the potential for errors, the requirement of significant time and resources, the need for specialist maintenance and updates, and the narrow scope of their application.³³ These appear to be some of the reasons for the ongoing limited take-up of these systems by legal and tax professionals.

2.2.2 Expert systems and current generative AI

Yang and Zhu further highlight that ‘[t]he current expert systems are not fully developed in integrating with emerging technologies. With the application of deep learning

²² Tan, above n 19, 2; Yang and Zhu, above n 16, 88559.

²³ Brock, above n 12, 9. He also points out that this process is recognised as having limits.

²⁴ Tan, above n 19, 2.

²⁵ Luger, *Artificial Intelligence*, above n 11, 20.

²⁶ Yang and Zhu, above n 16, 88559.

²⁷ Ibid 88561.

²⁸ William P Wagner, ‘Trends in Expert System Development: A Longitudinal Content Analysis of Over Thirty Years of Expert System Case Studies’ (2017) 76 *Expert Systems with Applications* 85, 85.

²⁹ Ibid 87.

³⁰ Durkin, above n 16, 8; Brock, above n 12, 8.

³¹ Durkin, above n 16, 11.

³² Yang and Zhu, above n 16, 88558.

³³ Ibid 88563.

technologies to expert systems, while enhancing the decision-making capabilities, the interpretability of expert systems has decreased'.³⁴

Legal expert systems (including tax systems) require a consideration of not only the legislation but also the various court decisions in the area and guidance from the taxation authorities. They require each to be broken down to a series of rules. In some situations, they require weightings to be attributed to different answers, particularly where more than one answer is available for each question.

2.3 Prolog

Prolog is one of the declarative languages based on predicate calculus. The name stands for programming in logic³⁵ and it is based on Horn clauses.³⁶ The difference between Prolog and the command-based languages is simply described. In Prolog the programmer tells the computer *how* to do something and in the command-based languages the programmer tells the computer *what* to do.³⁷ Prolog is based on facts and rules and the relationship between objects; the rules are often described as a database.³⁸

Rules in Prolog are expressed in the following manner:³⁹

friends(X,Y):- likes(X,Z), likes(Y,Z).⁴⁰

This may be interpreted as X and Y 'are friends if there exists a Z such that X likes Z and Y likes Z'.⁴¹ The facts are expressed in a like manner:⁴²

likes(george,kate).

likes(george,susie).

likes(susie,wine).

likes(george,wine).

Applying the foregoing rule, as Susie (X) and George (Y) both like wine (Z), on these facts they are friends.⁴³ If in the first two facts the order of the names was reversed then Kate and Susie would be friends as they both like George.

³⁴ Ibid 88581 (references omitted).

³⁵ Luger, *Artificial Intelligence*, above n 11, 538.

³⁶ I Gupta and G Nagpal, *Artificial Intelligence and Expert Systems* (Mercury Learning, 2020). A Horn clause consists of a result at its head and a set of conditions in its body.

³⁷ Ibid 359.

³⁸ Ibid 362.

³⁹ In Prolog variables use upper case; otherwise lower case must be used.

⁴⁰ The ':-' represents 'if', a comma represents 'and' and a semi-colon represents 'or'. An underscore in such a rule indicates no value is attributed to that item.

⁴¹ Luger, *Artificial Intelligence*, above n 11, 606.

⁴² In each example the first-named person likes the second-named person or item.

⁴³ Luger, *Artificial Intelligence*, above n 11, 604-608 gives a fuller description. The foregoing can be tested at <<https://swish.swi-prolog.org>>.

3. LARGE LANGUAGE MODELS

3.1 Introduction

The following describes LLMs succinctly:

LLMs are based on a transformer architecture and are characterized by their large size with billions of parameters. They have shown a comprehensive language understanding and the ability to produce text that is difficult for humans to distinguish from authentic human text in a wide range of language tasks.⁴⁴

As to their use in the professions that are centred on detail, a historian recently suggested:

[M]achine learning on enormous data sets is a science of particulars, more like, say, history than physics. With its data-driven, rather than reductionist, focus, this form of artificial intelligence threatens precisely the professions centered on particularity, on granular empirical knowledge, clinical knowledge, legal knowledge – and potentially historical knowledge.⁴⁵

Notwithstanding the promise and threat to some professions, on some occasions the output of an LLM may be meaningless. Whilst not a true word salad, something is wrong or missing from the logic; it is hallucinating. The LLM does not understand what it produces; the LLM simply applies probabilistic computations to create strings of text.

3.2 Some more history

As highlighted above, the current LLMs owe their development to the evolution of the transformer architecture.⁴⁶ Prior to that, work on natural language processing utilised a number of models, including statistical models (eg, n-gram), various forms of neural networks (including feedforward, recurrent, long short-term memory and gated recurrent units), knowledge-based models and contextual language models.⁴⁷ The other aspect to highlight in respect of LLMs is that, unlike the procedural languages and descriptive languages, LLMs are based, as already highlighted, on a probabilistic relationship between words.

The release of the GPT models and conversational AI ChatGPT in 2022 caught the attention of the general public. It is suggested that at the time of preparation of this

⁴⁴ Nils Constantin Hellwig, Jakob Fehle and Christian Wolff, 'Exploring Large Language Models for the Generation of Synthetic Training Samples for Aspect-Based Sentiment Analysis in Low Resource Settings' (2025) 261 *Expert Systems with Applications* 125514, 1 (references omitted).

⁴⁵ Matthew L Jones, 'AI in History' (2023) 128(3) *American Historical Review* 1360, 1361.

⁴⁶ Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser and Illia Polosukhin, 'Attention Is All You Need' (Conference Paper, 31st Conference on Neural Information Processing Systems, 4-9 December 2017). The transformer model is itself a type of neural network. See Aymen El Amri, *LLM Prompt Engineering for Developers: The Art and Science of Unlocking LLMs' True Potential* (FAUN, 2023) 37.

⁴⁷ El Amri, above n 46, 26; Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie and Ji-Rong Wen, 'A Survey of Large Language Models' (2024) *arXiv:2303.18223* (accessed 29 December 2024).

article there were over 5 million people building upon LLMs. They were doing so on platforms like OpenAI, Anthropic, Nvidia and Hugging Face. It was also suggested that the potential of this generation of AI models goes beyond simply natural language processing. It may be used, with varying success, in explaining complex algorithms or academic concepts, building interactive bots, and helping with software development.⁴⁸

Some of the open-source models available were Mistral, Llama and Gemma, and the closed-source models were OpenAI's ChatGPT models,⁴⁹ Cohere and Anthropic. The open-source models were available on platforms like Hugging Face. The closed-source models were available using an API⁵⁰ or often through other applications that access them. Access to those models was also available through various consumer products such as Microsoft Word (CoPilot) and more recently as part of the Apple operating system (Apple Intelligence).

As Bouchard and Peters highlighted, even the most powerful LLMs, such as GPT-4, rarely had domain-specific knowledge and often struggled with large volumes of data. All models may generate irrelevant and unreliable responses, or hallucinations.⁵¹ Such hallucinations have attracted attention in the legal industry because of the use of LLM outputs that were not carefully reviewed before their use.⁵² These irrelevant or unreliable responses may be controlled in many cases by changing various settings, as described below, tailoring the prompts, fine tuning and the use of domain-specific knowledge through the use of RAG files. There are occasionally advantages of allowing LLMs to hallucinate, one of which is that it may find a relevant novel solution.

3.3 Some LLM concepts

Reference has been made to current LLMs owing their development to the evolution of the transformer architecture and that they are based on a probabilistic model. At its simplest, a transformer examines multiple words (though sometimes the words are split, even down to letters in some cases) at once and understands the relationship between those words on a probabilistic basis.

A transformer consists of two main mechanisms, an encoder and a decoder. Both the encoder and decoder rely on neural networks and multi-head self-attention mechanisms.⁵³ The multi-head self-attention mechanism facilitates the dynamic weighting of the different words in the input. In doing so the model captures both the relationships and dependencies between the words.⁵⁴ As Ronald highlights, '[a]t its core, self-attention computes attention scores between each pair of words in the input

⁴⁸ Louis-François Bouchard and Louie Peters, *Building LLMs for Production: Enhancing LLM Abilities and Reliability with Prompting, Fine-Tuning, and RAG* (Towards AI, 2024) 662.

⁴⁹ GPT is an acronym for Generative Pre-trained Transformer. See El Amri, above n 46, 37.

⁵⁰ Bouchard and Peters, above n 48, 25, 662. An API is an application programming interface. As that expression implies, it is a piece of code that allows the program to access the functionality of another computer system, such as ChatGPT. See El Amri, above n 46, 37.

⁵¹ Bouchard and Peters, above n 48, 25, 662. Also see Nestor Maslej, Loredana Fattorini, Raymond Perrault, Vanessa Parli, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Juan Carlos Niebles, Yoav Shoham, Russell Wald and Jack Clark, *The AI Index 2024 Annual Report* (Institute for Human-Centered AI, Stanford University, 2024) 90 <https://hai.stanford.edu/assets/files/hai_ai-index-report-2024-smaller2.pdf>.

⁵² Dahl et al, above n 5,

⁵³ George J Ronald, *Mastering NLP with Transformers: Practical Techniques for Language Understanding and Generation* (Kindle edition, 2023) [2.2].

⁵⁴ Ibid.

sequence. These attention scores determine how much each word should attend to other words in the sequence when computing its representation. Higher attention scores indicate greater relevance between words'.⁵⁵

The result is a matrix of attention scores, with each element representing the relevance between a query word and a key word. They are then normalised and those 'attention weights are used to compute a weighted sum of the input representations',⁵⁶ resulting in the context vector. This context vector captures the aggregated information from all words in the sequence, weighted by their respective attention weights.

A vector allows for a representation of the text in a numerical form that the model can use. These representations are described as embeddings, which are high-dimensional vectors that capture the semantic relationship.⁵⁷ Tokens are the linguistic unit; they are the numerical representation of the text. As a general estimate, four characters or approximately 0.75 words of English text constitutes a token.⁵⁸

The total number of tokens that may be processed by an LLM model, both input and output, is constrained by the total number of tokens the model may handle.⁵⁹ The number of tokens that may be used in any situation may be further constrained by the user or a particular app.⁶⁰ Token considerations are important in interacting with LLMs, as they can constrain the interaction and its effectiveness.

A number of other parameters relevant in the use of LLMs⁶¹ are temperature, top-p and top-k.⁶² The temperature controls the variability in the output. A lower temperature provides a more focused output, whilst a higher temperature relaxes the constraints on the output and leads to hallucination and sometimes new insights.⁶³ The 'top-p' also controls the diversity in the output: 'It allows for the selection of the most probable words while limiting the selection to a subset of words whose cumulative probability exceeds a certain threshold'.⁶⁴

The 'top-k' parameter is used to control the number of possible choices when selecting the next words in a generated sequence. It limits the pool to the k most probable tokens at each generation step. This parameter enables the user to limit the diversity of the output.

3.4 Retrieval-augmented generation

RAG is another method that allows the user to enhance and improve the accuracy of an LLM response and reduce the likelihood of hallucinations. It improves the answer from an LLM by adding an information retrieval step from an external source before

⁵⁵ Ibid [2.3].

⁵⁶ Ibid.

⁵⁷ Janakiram MSV, 'The Building Blocks of LLMs: Vectors, Tokens and Embeddings', *The New Stack* (8 February 2024) <<https://thenewstack.io/the-building-blocks-of-llms-vectors-tokens-and-embeddings/>> (accessed 27 December 2024).

⁵⁸ El Amri, above n 46, [4.5].

⁵⁹ Ibid. As El Amri notes, the maximum number of tokens permitted by GPT-4 is 8192 and by GPT-3.5 turbo is 4096.

⁶⁰ This is often driven by cost considerations where closed models are used.

⁶¹ They may be included in API calls or set in models running locally.

⁶² El Amri, above n 46, [4.9].

⁶³ Ibid [4.9.1].

⁶⁴ Ibid [4.9.2].

providing the answer. Thus, it supplements the LLM's knowledge by adding relevant data from an external data repository into the LLM's input prompt. The source of data for a RAG may be internal documents, local databases and local code, or a series of published material.

The advantages of using a RAG include the ability to use a tailored source and domain-specific information, the ability to cite the source of the information it is relying on for its answer as part of the answer, and the ability for the internal data sources to be regularly updated to improve the RAG output.⁶⁵ Other methods of improving the output from LLMs include fine tuning⁶⁶ and prompt engineering. Using these methods in combination with a RAG is regarded as essential in adapting LLMs for more general use.

Figures 1 and 2 from LangChain depict the two separate stages in the RAG process.⁶⁷ The first is creating and loading the data by indexing the local repository.

Fig. 1: The RAG File Process




The second stage is using the indexed data by retrieving it from the data store as part of the prompt process and feeding it to the LLM before the LLM provides the answer.

⁶⁵ Bouchard and Peters, above n 48, 282-283.

⁶⁶ Ibid 663-664. This article does not discuss fine tuning.

⁶⁷ LangChain, 'Build a Retrieval Augmented Generation (RAG) App: Part 1' (Tutorial) <<https://python.langchain.com/docs/tutorials/rag/>> (accessed 30 December 2024).

Fig. 2: Use of RAG File Embeddings and Store

As the foregoing highlights, RAGs add a degree of complexity as they add integration complexity, often require data management of the RAG repositories and require accurate embeddings to achieve retrieval accuracy.⁶⁸ There are additional issues or risks with the use of RAGs. It is necessary to ensure the data repository is managed, is up to date and updated correctly, and the vector stores accurately reflect that data.⁶⁹ It is also necessary to ensure the use of the data in the RAG complies with ethical and regulatory requirements.⁷⁰ The use of a firm-based repository of documents may fulfil aspects of those requirements.

3.5 Prompt engineering

Prompt engineering is generally regarded as the simplest way to improve an LLM's performance for specific tasks. It is a quick and adaptable method for guiding the responses with specific inputs.⁷¹ Prompting has been described as vital in this area: 'Mastering the art of crafting *effective* prompts significantly enhances the performance of LLMs without requiring that models undergo underlying improvements'.⁷² Prompt engineering is itself an emerging discipline in the view of some in the LLM field,⁷³ whilst more recently it has been suggested that it will soon be overtaken by the ability of new LLMs.⁷⁴

⁶⁸ Bouchard and Peters, above n 48, 392.

⁶⁹ Ibid 411-412. There are various technical issues with their maintenance discussed at 411-420.

⁷⁰ Ibid 123, 412.

⁷¹ Ibid 391.

⁷² Maslej et al, above n 51, 148-149.

⁷³ El Amri, above n 46, [3].

⁷⁴ Dina Genkina, 'AI Prompt Engineering Is Dead', *IEEE Spectrum* (6 March 2024) <<https://spectrum.ieee.org/prompt-engineering-is-dead>> (accessed 3 January 2025).

By using well-formed and relevant prompts, adjusting the instructions and model parameters, and understanding the underlying model, more consistent and reliable outputs may be achieved.⁷⁵ Some simple steps to improve prompts include starting the prompt with an action verb rather than generic phrases, and providing a clear context in the question by including relevant details, underlying information and specific requirements. The prompt should be specific. It may assist to use role playing by asking the LLM to respond as the professional in the field. Double quotes can be used to emphasise particular words or phrases in the prompt. Text separators for different sections of the prompt may assist the LLM in understanding the structure of the prompt. Examples of the style of the output to be followed will also assist the LLM, as does an indication of the number of words to be included in the output. Other practices that may assist are guiding the model, refining the prompt, trying the prompt in different models, and experimenting and iterating.⁷⁶ OpenAI provides a list of example prompts for use with its models.⁷⁷

It should be remembered when using LLMs that they are probabilistic models trained on vast sets of data and do not understand the data, logic or maths. They just recognise patterns. Where the problems involve symbolic, mathematical or complex reasoning, it often assists to break the problem into steps, or to provide the LLM with examples that are easier to understand to improve the output.⁷⁸

The following is a very brief description of a few methods available to assist in improving prompts. ‘Few shot learning’ allows the LLM to make predictions based on a few examples.⁷⁹ ‘Chain of thought reasoning’ allows the user to provide reasoning for the answer. It may be used in combination with few shot learning⁸⁰ or on occasions with ‘zero shot prompting’.⁸¹ ‘Auto chain of thought prompting’ breaks the problem down into steps. One method is to indicate to the LLM by a suitable prompt that it will be prompted in a stepwise manner. Another method demonstrates to the model the questions and the step-by-step reasoning to be adopted. Both may be described as teaching the model by example.⁸² Another approach that is being developed is to use an LLM to iteratively generate prompts by breaking down the query.⁸³

Self-consistency is a further strategy that can enhance the model and the quality of the output. This involves exploring various paths to the answer using the LLM and can provide cross-checks to ensure the reliability of the output. A recent paper describes a different approach using soft logic (Horn clauses) in querying LLMs. The process described involves starting with a succinct distinct task query and an algorithmic automated stepwise process for repeatedly posing queries to an LLM that explores

⁷⁵ El Amri, above n 46, [3].

⁷⁶ Ibid [18].

⁷⁷ ‘Prompt Examples’, *OpenAI Platform* <<https://platform.openai.com/docs/examples>> (accessed 27 December 2024). See also El Amri, above n 46, [4.10]. There is also an open-source programming tool to test the perplexity of a prompt known as Betterprompt: Stylianos Iordanis, ‘betterprompt’, *GitHub* (2022) <<https://github.com/stjordanis/betterprompt>>. Another is promptfoo: ‘promptfoo’, *GitHub* (2025) <<https://github.com/promptfoo/promptfoo>>.

⁷⁸ El Amri, above n 46, [7].

⁷⁹ Ibid [6.1].

⁸⁰ Ibid [7].

⁸¹ Ibid [8]. ‘Zero shot learning’ means that there have been no prior examples or training on that task. It relies solely on the model’s pre-existing knowledge to generate the output.

⁸² Ibid [9].

⁸³ Maslej et al, above n 51, 150.

alternatives ('or') and adds detail ('and') to a fixed-length process. The process steers the exchange by updating the query with the responses using the LLM. At the completion there is a result and a pattern of reasoning.⁸⁴ So, using the model to assist in creating the prompts, because the model has familiarity with the concepts, enhances the quality of the output, lowering the perplexity. 'Transfer learning' is yet a further method. This method involves enhancing the learning of new problems by applying knowledge from a learnt method.⁸⁵ Avoid surprising the LLM, as the lower the prompt's perplexity, usually the better the output.⁸⁶

As the foregoing highlights, prompt engineering may be taken further with the use of stored queries and query engines, including sub-query engines. They aid the user by enhancing functionality, meeting the complexity of particular situations, and ensuring appropriate and adequate information is provided as a component of the query.⁸⁷ They often build in a number of the developing techniques, described above, that improve the output.

3.6 Local proprietary apps

Using a proprietary application may also assist in the development of a RAG-enhanced model and sometimes the prompt engineering. Those applications often allow for the indexing of the RAG repository, access to various LLM models both public and proprietary with appropriate API keys, and accessing LLM models running locally. The capabilities of these apps vary significantly.⁸⁸

3.7 Some legal LLM offerings: Australia

The three large legal publishing companies in Australia, Thomson Reuters, LexisNexis and Wolters Kluwer, now offer LLMs that include conversational search tools built on their repositories of reports, texts and other source material. These tools appear to be RAGs built on their vast repositories which are further refined by methods that include fine tuning, the control of the various parameters described, and reference citing, extraction and linking.

Thomson Reuters and LexisNexis are also offering tools that will allow the integration of a firm's or practitioner's documents into the LLM tool with the ability to ask questions of the documents, summarise them, create timelines and provide a comparison of the documents. The Thomson Reuters Australia website described its product⁸⁹ in these terms: 'Pose a question and scan through documents in minutes. CoCounsel delivers precise answers with source citations side-by-side for fast, efficient accuracy

⁸⁴ Paul Tarau, 'Full Automation of Goal-Driven LLM Dialog Threads with And-Or Recursors and Refiner Oracles' (2023) *arXiv:2306.14077*.

⁸⁵ El Amri, above n 46, [11].

⁸⁶ Ibid [12.1].

⁸⁷ Bouchard and Peters, above n 48, 397-398.

⁸⁸ In preparing this article using an Apple Silicon Mac, two such applications have been used. One of those applications provides its own server capabilities (Application B) and the other application (Application A) required access to a further application (LM Studio) that ran various downloaded models as the server. The further application was not used on a standalone basis, though it has the capability to be used in that manner. The applications and models used are not named as this is not considered relevant in the context of exploring the use of LLMs in a tax practice.

⁸⁹ Called CoCounsel. This was the description at the time of preparation of this article.

checks, enabling legal professionals to quickly uncover everything from crucial testimonies in extensive transcripts to important terms in complex contracts'.⁹⁰

A few providers offer document review, chat with your documents, drafting document assistance, multi-document review and research assistant facilities trained on Australian law. One such offering is styled AI Legal Assistant.⁹¹ Also some of the legal practice management software now incorporates LLM tools that provide assistance with legal research, assist in drafting documents with prompts and chat with the documents relating to a matter including the correspondence.⁹²

4. APPLICATION TO TAX PRACTICE

4.1 Introduction

The foregoing has provided an introduction to LLMs, some aspects of their use and attributes, RAGs, prompts in the use of LLMs, expert systems, decision trees, and Prolog and its role in creating expert systems. This article will now endeavour to bring that together by describing a study that used those methodologies to ask selected LLMs to provide advice on three different Australian taxation problems, two State and one federal. The study involved putting three different forms of query to LLMs. The first was a relatively simply query, in effect the baseline query. It put the question and facts simply to each model, unaided by an explanation of the rules or the relevant tax rates, where required. The same query was then put to the LLMs with a prompt that included the rules (and tax rates where required), and the third query had the rules (and tax rates where required) in a file (the RAG methodology) that was indexed and used by the LLMs in the process of considering the query.

The approach that was adopted of using Prolog-like rules, in the prompts and RAGs, was along the lines suggested by Eugene Asahara in *Enterprise Intelligence*⁹³ and his online blogs.⁹⁴ Other approaches are to use expert systems in Prolog that use an API call from Prolog to an LLM to improve the output.⁹⁵ A different approach is to use an LLM to extract from a natural language query the relevant information in the form of an attribute-value, feed that information into a Prolog programmed expert system, and then feed that result back into an LLM to create a natural language response.⁹⁶ A yet further approach is to integrate LLMs with structured knowledge bases to improve querying and output.

⁹⁰ Thomson Reuters, 'CoCounsel: Most Advanced Legal Assistant' <<https://www.thomsonreuters.com.au/en-au/products/cocounsel.html>> (accessed 24 December 2024).

⁹¹ AI Legal Assistant, above n 15. There are significant international commercial offerings, such as Harvey and Legora.

⁹² Leap, 'Exclusive AI: Purpose-Built for Law Firms' <<https://www.leap.com.au/legal-ai/>> (accessed 19 January 2025).

⁹³ Eugene Asahara, *Enterprise Intelligence: Bridging the Gaps between Wisdom, Business Intelligence, Knowledge Graphs, and Artificial Intelligence* (Technics Publications, 2024).

⁹⁴ Eugene Asahara, 'Soft Coded Logic' (Blog) <<https://eugeneasahara.com>>.

⁹⁵ See Richard de Rozario, 'Prolog to GPT API', *SWI Prolog* (February 2023) <<https://swi-prolog.discourse.group/t/prolog-to-gpt-api/6280>> (accessed 27 December 2024).

⁹⁶ Sanskar Sehgal and Yanhong A Liu, 'Logical Lease Litigation: Prolog and LLMs for Rental Law Compliance in New York' in Pedro Cabalar, Francesco Fabiano, Martin Gebser, Gopal Gupta and Theresa Swift (eds), *Proceedings 40th International Conference on Logic Programming (ICLP, 2024)* 59 <<https://cgi.cse.unsw.edu.au/~eptcs/paper.cgi?ICLP2024.4.pdf>>.

4.2 Methodology

As foreshadowed, three examples of tax problems were described in natural language and put to LLMs. In each case the problem was relatively simple and could form a base for a more complex problem by breaking down the rules into further rules. In each case five LLMs were used and one with two different temperatures, resulting in six outputs per type of query for each problem question run. Each time, the LLM was asked to write a memorandum of advice in response to the question relating to the particular tax situation described.

Each problem was put to the models using a particular application (described as Application A and Application B) in three stages. As already mentioned, the first stage involved the query and the facts being put to each model (a one shot query), the same query was again put to the models but with a prompt including the rules (and tax rates where required) in Prolog-like form, and the third had the same initial query (the one shot query) with the rules (and tax rates where required) in a RAG file that was indexed (tokenised) by the application and used by the application.

As far as possible the same models were used across the applications based on their description in the application and in the case of the fifth model using two temperatures: one was 0.2 and the other 1.5.⁹⁷ No adjustments were made to the other common parameters (top-p, etc) where permitted by the application.⁹⁸ Each request was set to 10,000 token use.⁹⁹ Similar base prompts were as far as possible used in all questions. Application A allowed the RAG embeddings to be encoded using one of the proprietary models. Application B allowed RAG embeddings to be encoded but it is unknown what embeddings were used. One law publisher model was used. However, it had a character limit and then no ability to use a RAG file. So, it was simply used on the first stage of each problem.

Table 1: Application and Model Details

Application	Model	Designation
A	Proprietary model, an older model	AM01
A	Proprietary model, a more recent model	AM02
A	Proprietary model, apparently a very recent model	AM03
A	Different recent proprietary model	AM04

⁹⁷ A temperature of zero requests no hallucination and 1.5 allows for a reasonable degree of hallucination. Application A was run with a temperature of 0.3 for each of the other models and Application B did not allow a temperature adjustment for the other models.

⁹⁸ Application A did not allow top-p or top-k parameters to be set. LM Studio did allow those parameters to be set as did Application B. Where they could be set they were left at the default settings of top-k 40, top-p 0.95 and min-p 0.05.

⁹⁹ This exceeds the tokens that may be processed by some of the models, as described earlier (n 59 above). It was anticipated that the application would throttle the use of any excess tokens, or the model would ignore the excess.

Application	Model	Designation
A	Publicly available model run locally with a temperature of 0.2	AM05
A	Publicly available model run locally with a temperature of 1.5	AM06
B	Proprietary model, an older model as per AM01	BM01
B	Proprietary model, a more recent model as per AM02	BM02
B	Proprietary model, apparently very recent. It is unclear if it is as recent as AM03.	BM03
B	Different recent proprietary model and not as recent as AM04	BM04
B	Publicly available model run locally with a temperature of 0.2 with similar base to AM05 but AM05 may have been further enhanced	BM05
B	Publicly available model run locally with a temperature of 1.5 with similar base to AM06 but AM06 may have been further enhanced	BM06

4.3 Problem 1: stamp duty

The *Stamp Duties Act 1923* (SA) (SDA) does not impose stamp duty on conveyances of qualifying land. All land is qualifying land other than land used for residential and primary production purposes, as defined in the SDA.¹⁰⁰ There are three criteria to determine whether land is used for residential purposes¹⁰¹ and two criteria for whether land is used for primary production.

There is a definition of the business of primary production but no definition of primary production. The section has been considered by the courts in South Australia in the context of whether land used for student accommodation constituted land used for residential purposes.¹⁰² Otherwise, the relevant provisions have not been considered by the courts and there are no similar interstate provisions.

The question and facts put to the LLMs in each case were as follows:

¹⁰⁰ *Stamp Duties Act 1923* (SA) ('SDA') ss 105A(1), (5).

¹⁰¹ *Ibid* s 105A.

¹⁰² *Perpetual Corporate Trust Ltd v Commissioner of State Taxation* [2022] SASC 7; *Commissioner of State Taxation v Perpetual Corporate Trust Ltd* [2022] SASCA 117.

Stamp duty in South Australia is not imposed on qualifying land.^[103] Please provide a memorandum of advice for a client advising the stamp duty consequences of the conveyance of the following pieces of land using the following prolog like rules by reference to the following identifiers of the land. The structure of the memorandum should be along the following lines. Set out the facts, describe the likely classification of the respective holding and describe the tax consequences of acquiring the land with reasons and computation of any tax consequences of doing so. Use as many words as required to provide all of information requested.

#####

%land(Identifier, Use, Improvements, Zoning, Valuer Generals Primary Production Determination, Capital Value)

land(blackacre, industrial, no_improvements, no_determination, no_determination, 750000).

land(greenacre, residential, residential_character, predominant_use, predominant_use, 2000000).

land(blueacre, commercial, commercial, residential_character, none, 3000000).¹⁰⁴

land(greyacre, residential, minor, residential_zoning, none, 2000000).

land(yellowacre, primary_production, minor, residential_zoning, none, 4000000).

land(brownacre, no_use, no_improvements, primay_production_zoning, primary_production, 350000).¹⁰⁵

The Prolog-like rules (including duty rates) which were added to the prompt on the second round and included in the RAG file in the third round are set out in Appendix 1, Part C.

It may be expected that the answer would conclude that blackacre and blueacre both satisfy the qualifying land requirement and could be conveyed stamp duty free. The conveyance of the remaining land would attract duty. The duty computations then required an application of the stamp duty rates in Part 1 of Schedule 2, Item 3 of the SDA to the market values provided. There was no guidance provided in the prompt as to whether any particular criteria prevail, save for possibly the order of the criteria or from an understanding of the legislation.

¹⁰³ Note that this is a statement of fact but not further explained. It is left for the model to understand.

¹⁰⁴ This fact situation is apt to create an issue as both the use and improvements are commercial but the zoning is residential.

¹⁰⁵ As described earlier, the underscore in such rules indicates no value is attributed to that item.

4.3.1 Summary of outcomes

The outcome of using each of the LLMs in the manner described is summarised in each of the Tables in Appendix 1 Part A. Appendix 1 Part B provides two examples of the outputs, one using a single shot prompt and the other using Prolog-like rules in the prompt. None of the models using the single shot prompt produced correct or useful output. Some of the responses were nonsensical. The layout and approach adopted by model AM03 was both useful and helpful and could with the necessary corrections be used as a basis for providing the appropriate advice.

Two of the models using Prolog-like rules and rates in the prompts produced meaningful and useful outputs. Both of these outputs could with some minor modifications be used to provide the substantive advice required. The other outputs could not be realistically used. The two useful models were then recent LLMs, highlighting the improvement in the LLMs available. This suggests that the use of well-enhanced prompts from a prompt library with carefully selected models will produce useful results.

The models using a RAG file with Prolog-like rules were disappointing. Two were similar, but otherwise, the output from the models using the Prolog-like rules in the prompt were often better. This was unexpected and is discussed below. It is possible that the RAG approach will be improved by running many examples and fine tuning the model.¹⁰⁶

Table 2 summarises the usefulness of the outputs of the models across the three different processes. The Table splits the outputs in each case into three classes: those that were useful, namely they accurately addressed the issue and could be used as a good basis to provide the required advice (an 'X' is used to indicate those models), those that were close and only required some modifications or corrections to be able to be used to provide the advice (an 'x' is used to indicate those models), and those of limited use. Some would be of use with considerable work and some were so confusing that they would be better discarded. Tables 5 to 7 (Appendix 1 Part A) review each output in much more detail.

Table 2: Stamp Duty Example – Summary of Outcomes

Model	One shot		Full prompt		RAG file	
	Useful	Limited	Useful	Limited	Useful	Limited
AM01		x		x		x
AM02		x	x			x
AM03		x	X		x	
AM04		x	X		x	
AM05		x		x		x

¹⁰⁶ See n 66 above.

Model	One shot		Full prompt		RAG file	
	Useful	Limited	Useful	Limited	Useful	Limited
AM06		x		x		x
BM01		x		x		x
BM02		x		x		x
BM03		x	X			x
BM04		x	X		x	
BM05		x		x		x
BM06		x		x		x
CM07		x				

4.4 Problem 2: land tax

Land tax in South Australia is imposed by the *Land Tax Act 1936* (SA) (LTA) based on the site value of the land. It is not imposed on a person's principal place of residence,¹⁰⁷ primary production land outside the defined rural area¹⁰⁸ and in a number of other situations that are not relevant for the purpose of this example.¹⁰⁹

Aggregation also applies, so if a taxpayer owns more than one piece of land then those properties are aggregated to determine the land tax payable.¹¹⁰ Certain trust holdings are taxed at a different rate to individuals and companies.¹¹¹ Companies may be grouped and treated as a single owner in certain situations.¹¹²

The question put to the LLMs was as follows:

Jim is married to Mary. They have two adult children John and Judy. Jim and Mary control three family trusts. The beneficiaries of those family trusts include Jim, Mary, their children and remoter issue and any companies or trusts in which they have an interest. Interest includes in the case of a company as a shareholder or are owed money and in the case of a trust are a beneficiary or are owed money (called a Relevant Trust). Jim and his family are also interested in the company if a Relevant Trust is a shareholder.

¹⁰⁷ *Land Tax Act 1936* (SA) ('LTA') ss 5(10)(a)-(bb). There are many interstate decisions as to whether land is used as a principal place of residence.

¹⁰⁸ *Ibid* s 4(1)(m).

¹⁰⁹ *Ibid* ss 4, 5.

¹¹⁰ *Ibid* s 8B.

¹¹¹ *Ibid* s 11.

¹¹² *Ibid* Pt 3, Div 6.

Jim proposes to acquire a further property described below and is seeking advice as to the most suitable taxpayer (entity) from a land tax perspective only in which to acquire the property based on the following information.

Please provide a memorandum of advice describing the taxpayer (entity) including a new taxpayer (entity) to acquire the property with the least land tax ongoing cost having regard to the following information.

The structure of the memorandum should be along the following lines. Set out the facts, describe the current tax consequences of the respective holdings, describe the tax consequences of acquiring the new property by any of the existing entities, the recommendation as to the manner of acquisition with reasons and computation of the tax consequences of doing so.

#####

%land(Identifier, Use, Owner, Site_value)

existing_land(blackacre, commercial property, Jim, 2000000).

existing_land(whiteacre, commercial property, Mary, 1500000).

existing_land(brownacre, commercial property, Jims Family Trust, 3000000).

existing_land(greenacre, primary production, Marys Trust, 2500000).

existing_land(blueacre, principal residence, May,^[113] 1250000).

existing_land(yellowacre, residential, Marys Trust, 650000).

existing_land(cityblock, commercial properties, Jim Mary Trust, 4500000).

new_property(highville, commercial shopping centre, _, 3250000).

The Prolog-like rules (including the tax rates) which were added to the prompt on the second round or included in the RAG file in the third round are set out in Appendix 2 Part B.

It may be expected that the answer would conclude that highville could be acquired by a new discretionary trust, a new company or one of the children. The advantage of a company is that, based on the existing facts provided, a company is not currently used by Jim and his family and the tax would be at the standard rates. The property could

¹¹³ Whilst the problem could have been run again with this corrected, it was decided to leave it to highlight how the models treat such a simple typographical error. Most models did not treat May as Mary, notwithstanding the initial facts named the family members but contained no mention of May. May is described as owning the principal residence, which is exempt from land tax, so it added no value to Mary's holdings in any event (LTA, above n 107, s 5(10)(a)). Mary's other holding is a commercial property that would attract land tax.

also be acquired by one or both of the children and would be assessed at the standard rates.

In practice other taxation and control considerations usually result in the acquisition by a new discretionary trust with a corporate trustee. The tax computations would then require an application of the land tax rates. The question did not specify the financial year and therefore one would assume the current financial year.

The land tax rates for the current financial year are not set out in the LTA as the thresholds for their application are indexed and they change yearly.¹¹⁴ Therefore rates are usually ascertained from the Commissioner's website and publications. As also noted above, the facts included a misspelling of Mary's name, which was not corrected.

4.4.1 *Summary of outcomes*

The outcome from each of the LLMs was summarised in Tables similar to those in Appendix 1 Part A.¹¹⁵ Appendix 2 Part A provides an example of the output using the Prolog-like rules in the prompt. A number of the models using the single shot prompt produced a useful output. Some of the responses were nonsensical. The layout and approach adopted by some of them was useful and helpful. They could with the necessary corrections be used as a basis for providing the appropriate advice.

A number of the models using Prolog-like rules and rates in the prompts produced meaningful and useful outputs. AM04 was particularly useful. A number of those outputs could with some minor modifications be used to provide the substantive advice required. The useful models were again the recent LLMs, highlighting the improvement in the LLMs available. This again suggests that the use of well-enhanced prompts from a prompt library with carefully selected models also achieves a suitable output.

The models using a RAG file with Prolog-like rules were again generally disappointing (see the discussion in sections 4.3 and 5.3). The model AM03 would be particularly useful and a couple of the others would be useful with some modification.

Table 3 summarises the usefulness of the outputs of the models across the three different processes. The Table describes the outputs in the same manner as in Table 2.

Table 3: Land Tax Example – Summary of Outcomes

Model	One shot		Full prompt		RAG file	
	Useful	Limited	Useful	Limited	Useful	Limited
AM01	x		x			x
AM02	x			x		
AM03	x		x		X	
AM04	x		X		x	
AM05		x		x		x

¹¹⁴ LTA, above n 107, s 8A, Sch 1.

¹¹⁵ They are not reproduced in this article.

Model	One shot		Full prompt		RAG file	
	Useful	Limited	Useful	Limited	Useful	Limited
AM06		x		x		x
BM01	x		x			x
BM02	x		x		x	
BM03	x		x			x
BM04	x		x			x
BM05		x		x		x
BM06		x		x		x
CM07	x					

4.5 Problem 3: Division 152 – small business relief

Division 152 of the *Income Tax Assessment Act 1997* (Cth) provides for various forms of relief from capital gains tax for small businesses where a number of conditions are satisfied. There are four basic conditions,¹¹⁶ though some have internal alternatives, and each of those conditions involves a number of further requirements.

The four requirements are as follows:

- (a) a capital gains tax event happens in relation to a capital gains tax asset of the taxpayer in an income year;
- (b) the event would (apart from Division 152) have resulted in the gain;
- (c) at least one of the following applies:
 - (i) the taxpayer is a capital gains tax small business entity for the income year;
 - (ii) the taxpayer satisfies the maximum net asset value test;
 - (iii) the taxpayer is a partner in a partnership that is a CGT small business entity for the income year and the CGT asset is an interest in an asset of the partnership;
 - (iv) the conditions mentioned in subsection (1A) or (1B) of section 152-10 are satisfied in relation to the capital gains tax asset in the income year;

¹¹⁶ *Income Tax Assessment Act 1997* (Cth) s 152-10(1).

(d) the capital gains tax asset satisfies the active asset test.¹¹⁷

There are additional rules for interests in partnerships, passively held assets, shares in a company and interests in trusts.¹¹⁸ The satisfaction of the basic conditions allows a taxpayer to choose a 50 per cent reduction¹¹⁹ in the gain and an eligibility to choose rollover relief.¹²⁰ Complete relief from any capital gain is available if the 15-year exemption requirements are satisfied.¹²¹ There is also a small business retirement exemption subject to certain conditions,¹²² up to an individual's lifetime limit of AUD 500,000.¹²³ These two have additional rules.

The question put to the LLMs was as follows:

Jim has sold his business and is keen to know whether the basis conditions for Division 152 of the Income Tax Assessment Act 1997 (Cth) are satisfied. If they are and without further consideration he is entitled to an additional 50% capital gains discount. Please provide a memorandum of advice to be given to Jim describing whether he is entitled to the Division 152 relief. The structure of the memorandum of advice should describe the facts, whether Jim satisfies the basic conditions for the Division 152 and the basis for your decision with a conclusion.

#####

Division 152 Elements

(Entity, CGT Event, Gain, CurrentYearIncome, LastYearIncome, NextYearIncome, GrossAssets, Partner, Conditions1A1B, CGTAsset, AssetPartnership, ActiveAsset)

(Jim, a1, 1900000, 1200000, 2400000, 1000000, 8000000, no, no, no, no, yes)

The Prolog-like rules which were added to the prompt on the second round or included in the RAG file in the third round are set out in Appendix 3 Part B.

It may be expected that the answer would conclude that Jim satisfies the four basic conditions for Division 152 to apply, entitling him to the additional 50 per cent discount allowed by subdivision 152-C. That is on the basis that a CGT event in relation to a CGT asset occurs on the sale of the business, that is, it is CGT event A1 (which is provided in the facts). The facts indicate that there is a CGT gain of \$1,900,000. So the first two basic conditions are satisfied.

The third basic condition provides for a number of alternatives, any one of which satisfy the condition. The first is that the taxpayer is a small business entity. The facts disclose

¹¹⁷ Ibid.

¹¹⁸ Ibid s 152-10(1A), (1B), (2).

¹¹⁹ Ibid sub-div 152-C.

¹²⁰ Ibid sub-div 152-E.

¹²¹ Ibid sub-div 152-B.

¹²² Ibid sub-div 152-D.

¹²³ Ibid s 152-320.

that Jim is selling his business and describes his income in the two preceding years and the current year (a proxy for his turnover). On the facts provided Jim appears to satisfy the small business entity test. The second is that the taxpayer satisfies the maximum net asset test; the net asset test has a limit of \$6,000,000. The facts disclose Jim's net assets are \$8,000,000. Jim therefore fails that test. The facts also indicate that Jim is not in a partnership, nor do subsections (1A) or (1B) apply. On the basis of Jim being a small business entity, he satisfies the third basic condition. The fourth basic condition is that the asset is an active asset. Once again that is a stated fact and is therefore satisfied.¹²⁴

4.5.1 Summary of outcomes

The outcome from each of the LLMs was summarised in Tables similar to those in Appendix 1 Part A.¹²⁵ Appendix 3 Part A provides an example of the output using Prolog-like rules in the prompt. A few of the models using the single shot prompt produced a useful output and with corrections could be used as a basis for advice. Some of the responses were nonsensical.

A number of the models using Prolog-like rules and rates in the prompts produced meaningful and useful outputs. AM03 was particularly useful. A number of those outputs could with some minor modifications be used to provide the substantive advice required. Again, the useful models were recent LLMs, highlighting the improvement in the LLMs available. Again, this suggests that the use of well-enhanced prompts from a prompt library with carefully selected models also achieves a suitable output.

The models using a RAG file with Prolog-like rules were again generally disappointing. Model AM03 was particularly useful and one other would be useful with some modification. Once again, the outputs from the models using the Prolog-like rules in the prompt were often better (see the discussion in sections 4.3 and 5.3).

Table 4: Taxation Example – Summary of Outcomes

Model	One shot		Full prompt		RAG file	
	Useful	Limited	Useful	Limited	Useful	Limited
AM01		x		x		x
AM02	x		X			x
AM03	x		X		x	
AM04		x	x		X	
AM05		x		x		x
AM06		x		x		x
BM01		x		x		x

¹²⁴ As stated earlier, obviously these conditions can be further broken down into their various subrules.

¹²⁵ They are not reproduced in this article.

Model	One shot		Full prompt		RAG file	
	Useful	Limited	Useful	Limited	Useful	Limited
BM02		x	X			x
BM03		x		x		x
BM04		x	x			x
BM05		x		x		x
BM06		x		x		x
CM07	x					

5. CONCLUSION AND SOME IMPLICATIONS

5.1 Introduction

This article has explored the use of LLMs in preparing tax advice that describes the tax outcomes and/or tax consequences of a given fact situation, rather than the more general use of LLMs. The exploration was undertaken against the background of the long use and development of expert systems to provide an analysis of a given situation, often using Prolog. The article then turned to the use of Prolog-like rules as a possible method of enhancing the output from LLMs.

5.2 Discussion and synthesis of the results

In the case of each of the detailed Australian tax fact situations used, two State and one federal, the results highlight that the use of LLMs unaided to assist practitioners in providing detailed tax advice is of limited assistance and on some occasions unhelpful. It was expected that the federal tax query would produce a better output as it was anticipated that the LLM would have encountered more webpages and material discussing federal tax than State taxes, particularly South Australian State taxes. As Tables 2, 3 and 4 highlight, that did not occur.

The outcome at this point is not unexpected, as others have observed, the performance of LLMs ‘deteriorates when dealing with more complex tasks that require a nuanced understanding of legal issues or interpretation of legal texts’ and ‘different models exhibit varying degrees of accuracy and biases’.¹²⁶

Also, for reasons that are not particularly obvious,¹²⁷ save for possibly the more general nature of the question, the unaided prompt in respect of the land tax example provided generally better outcomes than the federal tax query. It was also apparent that the models used did not know the current tax rates, where they were required to aid the analysis,

¹²⁶ Dahl et al, above n 5.

¹²⁷ One of the usual criticisms of LLMs is that they are in effect ‘black boxes’ and they cannot obviously explain their reasoning: Adel Abusitta, Miles Q Li and Benjamin CM Fung, ‘Survey on Explainable AI: Techniques, Challenges and Open Issues’ (2024) 255(Pt C) *Expert Systems with Applications* 124710.

though sometimes the position may be inferred without knowing the rates, which some models demonstrated.

Where the prompt was enhanced with Prolog-like rules (and rates), the LLMs produced a number of useful outputs. Those outputs were well laid out and often included computations (where required) to support the analysis. Such computations were of considerable assistance in understanding the reasoning. They often provided a good basis for providing advice to a client. Occasionally there were minor errors that could be easily corrected. These models tended to be the newer proprietary models.¹²⁸ The output of the open-source models that allowed, through the applications used, the modification of the temperature (and some of the other parameters) were generally not useful.¹²⁹

Obviously enhancing the prompt with detailed rules and rates requires significant effort. Sometimes the models can be used to assist this process. It is suggested that such rules then need to be created as part of a suite of standard prompts (a prompt or snippet library) so that they can be reused and added to, where required, to create further prompts with some of the rules further broken down. By way of example, the division 152 question in this study simply used the four basic conditions. Each of those basic conditions could be broken down into a series of further rules both by reference to the legislative provisions and some decided cases. Also, where this process is adopted, the prompts should incorporate suitable guidance on the structure of the facts to be provided.

5.3 Some of the lessons learned

The more modern proprietary models produced much more useful output than the older models. This suggests that at least for the time being the focus should be on the use of those models. More models have been released since this study was undertaken. However, without the prompts being enhanced with the Prolog-like rules (and rates) the output of the newer models is only marginally better than the results reported in this article.

So, for the moment, the use of LLMs to provide detailed tax advice in respect of detailed fact situations unaided by extensive prompts or RAG files appears unrealistic, as Tables 2, 3 and 4 highlight. This means that tax professionals must put considerable effort into creating extensive prompt libraries or developing simpler prompts utilising RAG files with the rules and rates.

The law publishing company model had a character limit, so this model was not available for use in this manner. With further enhancements to that model and the removal of such restrictions, it is anticipated this model will be even better, being enhanced by the publisher's reservoir of legal material.

The RAG file-enhanced output, in each case, as described, was generally disappointing. The expectation, though with hindsight, possibly unrealistic, was that it would not be

¹²⁸ Whilst the newer models may be much better, their performance may change over time because of updating by the proprietors and in some cases may deteriorate over time. See Maslej et al, above n 51, 143.

¹²⁹ If the analysis was undertaken outside of those applications using one of the more common computer languages, such as Python, then API calls could have been made that included temperature and other control criteria with the enhanced prompt.

much different from the Prolog-like rules enhanced prompt output. As Tables 2, 3 and 4 highlight, that was not the case.

There appear to be several reasons why the RAG approach was disappointing. The following only touches on some of them. The first is that, with rules in the prompt, the LLM is required to consider all the rules. The structured, conditional, if-then form of Prolog-like rules essentially creates an implicit program, and the LLM applies it through something resembling chain-of-thought symbolic execution.

However, when the same rules are provided to the LLMs through RAGs several factors impact their effective use. LLMs rely on semantic search and vector space analysis; RAG simply enhances both. Further, simple or naïve RAG retrieves numerous fragmented context chunks, the chunks of text that have been vectorised, to be put into the LLM context window. This chunking may not only create fragmentation but also may cause loss of interdependence. The fragmentation leads to a lack of understanding and reduces the effectiveness of the RAG approach. So, the RAG approach, in this context, suffers from both rule fragmentation and semantic search problems.¹³⁰

Importantly, in the prompt the Prolog-like rules provide a complete set of logical constraints that must be observed. In a Prolog environment, if the rules are not satisfied, the query fails.¹³¹ In an LLM environment, the model will seek to solve the problem by both attempting to solve the question and using only some of the rules according to the required ranking. In the tax context the difference in approach can be significant. Generally, each of the positive rules must be satisfied; there may also be negative rules that must be satisfied and which may be semantically distant in the model, and there are often exceptions that may also be semantically distant. A detailed consideration of those reasons, and possibly others, could be the subject of a further, more focused technical article.¹³² In conclusion, good outputs from an enhanced prompt can provide considerable assistance in advising on the tax outcomes and/or tax consequences of a given fact situation. Further, once the prompt library is created, much like most precedents, it could be used to produce useful reasoned output that can quickly form the basis of advice.

The implications for tax practitioners are simple. LLMs are useful tools and, when suitably enhanced, they can be used to provide detailed reasoned tax advice. However, they require extensive preparatory work for day-to-day use, whether it is picking the current models or developing the prompt library. As the LLMs improve, those prompt libraries and possibly RAG files could be used with those improved LLMs to improve

¹³⁰ Keith Bourne, *Unlocking Data with Generative AI and RAG: Enhance Generative AI Systems by Integrating Internal Data with Large Language Models Using RAG* (Packt Publishing, Kindle Edition, 2024) chs 8 and 14, and 529.

¹³¹ Luger, *Artificial Intelligence*, above n 11, ch 2, 314-317, 542-543.

¹³² See Gary Marcus, 'The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence' (2020) *arXiv:2002.06177*; Gary Marcus, 'Deep Learning: A Critical Appraisal' (2018) *arXiv:1801.00631*; Subbarao Kambhampati, 'Can Large Language Models Reason and Plan?' (2024) 1534(1) *Annals of the New York Academy of Sciences* 15; Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-Tau Yih, Tim Rocktäschel, Sebastian Riedel and Douwe Kiela, 'Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks' in Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan and Hsuan-Tien Lin (eds), *Advances in Neural Information Processing Systems* 33 (Curran Associates, 2020) 9459; Maslej et al, above n 51, 116. I am also indebted to an anonymous reviewer for suggesting that this issue be given further consideration and Anthropic's Claude for some of its suggestions.

the output, though they are likely to require updating. So, whether developed in the form of prompt libraries, with Prolog-like rules and rates, or possibly RAG files with those rules and rates including a suitable library of more limited prompts, that work is likely to be able to be re-used for some years whilst the LLMs are further improved.

The challenges are not only to adopt these tools and develop the prompt libraries (much like current precedent libraries) but to educate tax practitioners on their effective use and more importantly to understand when the output, though looking impressive, is wrong. Whilst the more senior practitioner is likely to be reasonably adept at doing so, the younger practitioner may not be. This is a matter of ensuring that the appropriate process is in place to address this possibility.

As Dahl and co-authors observe, ‘integration of AI in legal practice will require more iteration, supervision, and human understanding of AI capabilities and limitations’.¹³³ This will obviously be a bigger challenge for smaller firms with limited resources than larger firms.¹³⁴ The availability of LLM enhanced tools will also likely lead to greater lawyer retention, as a United Kingdom survey indicated, 92 per cent of students are using LLMs.¹³⁵ Those students will expect such tools to be generally available and, if not, they will move.

6. APPENDICES

APPENDIX 1 PART A: STAMP DUTY EXAMPLE: SUMMARY OF OUTPUT

Table 5: Stamp Duty Example – Single Shot Prompt

LLM	Description of output
AM01	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption or calculate the duty where payable. It provided vague general advice including the following statement: ‘Seeking professional advice from a conveyancer or tax advisor would be beneficial to ensure compliance with stamp duty regulations.’

¹³³ Dahl et al, above n 5.

¹³⁴ It is anticipated the situation in Australia is much like the United Kingdom in this respect, where larger firms look positively to generative AI whilst mid-tier and smaller firms expect it to impact negatively on their profitability. PwC, *Annual Law Firms’ Survey 2024: Innovate, Adapt, Protect: A Blueprint for Sustained Growth* (2024) 6-7 <<https://www.pwc.co.uk/industries/legal-professional-business-support-services/law-firms-survey.html>>.

¹³⁵ Josh Freeman, *Student Generative AI Survey 2025* (Higher Education Policy Institute Policy Note 61, 2025) <<https://www.hepi.ac.uk/2025/02/26/student-generative-ai-survey-2025/>>.

LLM	Description of output
AM02	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption, or calculate the duty where payable. For some reason it stated that the primary production land 'is a qualifying land for exemption' so brownacre and yellowacre will be exempt. It also stated: 'Please note that the exact computation of stamp duty depends on the rates provided by RevenueSA. The rates may vary and are usually progressive, with the rate increasing as the property value increases.'
AM03	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption, or to correctly calculate the duty where payable. For some reason it stated that the primary production land is exempt and proceeded to incorrectly apply that proposition. It described stamp duty rates that were incorrect and incorrectly calculated any duty. It is not clear where the rates used came from. It was confused by the yellowacre fact situation. The layout of the answer with its facts, individual description and summary was both useful and helpful.
AM04	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption, or to calculate the duty where payable. For some reason it stated that the primary production land is exempt and proceeded to incorrectly apply that proposition. It was confused by the yellowacre fact situation and stated that required further investigation. The layout of the answer with the facts, individual description and summary was helpful. It had a final note that 'This advice is based on the provided rules and should be verified against current legislation and regulations.'
AM05	The model assumed that the recipient of the memo was acquiring the land. The explanation for the exemptions and their application was confusing. The model correctly identified blackacre as qualifying land but did not explain why. It was confused by both blueacre and yellowacre. It said blueacre and yellowacre are subject to stamp duty in its discussion but did not calculate any duty consequences.
AM06	The model assumed that the recipient of the memo was acquiring the land. It incorrectly described the SDA and the stamp duty rates including a statement that no stamp duty is payable on a capital value of less than \$750,000. It is not obvious where that rate came from. Blackacre and brownacre are both exempt on that basis. All other land is dutiable. It concluded with, inter alia: 'It is recommended that the client consults with a qualified real estate agent or solicitor to ensure they are aware of all the implications of acquiring these pieces of land. Note: This memorandum is for general purposes only and should not be relied upon as professional advice.'
BM01	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption or calculate the duty where payable. It stated that greenacre as residential land and brownacre as primary production land are exempt. There was no attempt to calculate any duty consequences. The model provided vague general advice including the

LLM	Description of output
	following statement: 'Please note that individual circumstances and specific details may impact the stamp duty consequences, and it is advisable to seek legal advice tailored to the specific situation to ensure compliance with all relevant laws and regulations.'
BM02	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption, or to calculate the duty where payable. For some reason it stated that the primary production land 'is a qualifying land for exemption' so brownacre is exempt but not yellowacre because the Valuer-General has not classified it as primary production land. It concluded with: 'We recommend consulting a taxing authority or legal advisor to obtain precise calculations and confirm the duty payable for each property.'
BM03	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption, or to calculate the duty where payable. For some reason it stated that the primary production land 'is typically exempt'. It concluded: 'Each parcel of land invokes distinct stamp duty considerations rooted in land use, zoning, and improvements as outlined by South Australian statute and guidelines. It is imperative to consult the most current legislative framework and possibly the South Australian Revenue Office or legal expertise for detailed duty computations and exemptions that could apply. This memorandum summarizes potential liability areas, yet advises further detailed investigation based on specific legislative updates and nuanced legislative interpretations to ensure compliance and potential duty abatements.'
BM04	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption, or to calculate the duty where payable. It assumed primary production and residential land is exempt. It concluded that brownacre is most likely to be exempt from stamp duty as qualifying land. Greenacre and greyacre are also likely to be exempt, but further confirmation may be needed. Yellowacre's status is uncertain and requires additional investigation. Blackacre and blueacre are likely to be subject to stamp duty as non-qualifying land. It concluded: 'It is recommended to seek further clarification from the relevant South Australian authorities regarding the exact criteria for qualifying land, especially for properties with mixed characteristics like Yellowacre. Additionally, up-to-date stamp duty rates should be obtained for accurate calculations on non-exempt properties.'
BM05	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption, or to calculate the duty where payable. It overlooked brownacre. It stated that greenacre is exempt because it is residential land. It concluded with a recommendation: 'you should consult with a stamp duty specialist or tax advisor to ensure compliance with all relevant laws and regulations.'
BM06	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption, or to calculate the duty where payable. It

LLM	Description of output
	misdescribed some of the land, blackacre being described as primary production land. It also said that the disposition will have a capital gains tax consequence. Greenacre and yellowacre will be dutiable as residential land as they ‘fall within a “zero” rate bracket’. It continued: ‘however no benefits apply if either property is later transferred to a commercial use or the other party uses it in such manner.’ The summary was then inconsistent with the analysis in some areas and stated: ‘blackacre and greyacre are likely eligible for an exemption from stamp duty due to primary production zoning; - greenacre benefits zero-rate status but may need re-assessment should use or intended land post transfer occur; - blueacre falls outside of the ‘residential’ bracket due to commercial property being a feature so does fall within Capital Gains Tax; - brownacre is not zoned residential (it has been noted it doesn’t appear eligible for either classification in terms stamp duty). Please note that this memo serves only as an advisory, and you are strongly advised to consult directly with the relevant authorities and/or professional tax advisor prior to any land transaction.’
CM07	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption, or to calculate the duty where payable. It correctly referenced s 105A of the SDA. It concluded that brownacre is exempt as primary production land. It misdescribed the rates on values below \$500,000.

Table 6: Stamp Duty Example – Prompt with Prolog-Like Rules

LLM	Description of output
AM01	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption or to calculate the duty where payable. It stated that greenacre, as residential land, was exempt, greyacre was also exempt as residential land and yellowacre was exempt as it was classified as primary production land with minor improvements and residential zoning. However, brownacre was not exempt because of its primary production zoning. Notwithstanding the Prolog-like rules with the rates and the capital values were provided, the model stated: ‘To calculate the stamp duty payable for each piece of land, we would need to determine the capital value and apply the relevant stamp duty rate based on the total amount.’
AM02	The model correctly classified the land other than blueacre as qualifying or non-qualifying land. In the case of blueacre the model said: ‘it may be considered non-qualifying land, and hence, stamp duty would be payable’. The language in each case was tentative. The model did not calculate the stamp duty payable in respect of each dutiable lot and the one example of the duty consequences it offered was incorrect. It concluded: ‘It is recommended to consult with a tax advisor or legal professional for a detailed analysis.’

LLM	Description of output
AM03	The model correctly classified the land other than blueacre as qualifying or non-qualifying land. It provided a stamp duty calculation for each piece of land it regarded as dutiable in a well laid out memo with a convenient summary of the tax consequences.
AM04	The model correctly classified the land and calculated the duty on the dutiable land. There were some minor computation errors. It was well laid out and much simpler than AM03.
AM05	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption. It exempted primary production land. Where it did calculate the stamp duty it was significantly incorrect.
AM06	The model correctly classified blackacre and blueacre as qualifying land though the explanation was flawed. The model incorrectly classified yellowacre and brownacre as exempt as primary production land. There was no attempt to calculate the duty and the concluding statement oddly said: 'The above calculations assume that the client has no other liabilities or obligations that may affect the tax consequences of acquiring these lands.' It further stated: 'It is recommended that the client consults with a tax professional to ensure compliance with all applicable laws and regulations.'
BM01	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption or to calculate the duty where payable. It stated that greenacre and greyacre were exempt as residential land and yellowacre and brownacre were exempt as primary production land. There was no attempt to calculate the duty.
BM02	The model correctly classified blackacre. It incorrectly classified yellowacre and brownacre as exempt as primary production land. Where it calculated duty it made errors.
BM03	The model correctly classified the land other than blueacre as qualifying or non-qualifying land. It provided a correct stamp duty calculation for each piece of land it regarded as dutiable. It was well laid out.
BM04	The model correctly classified the land other than blueacre as qualifying or non-qualifying land. It provided some incorrect stamp duty calculation for each piece of land it regarded as dutiable. It was well laid out.
BM05	The model failed to classify most of the land accurately for the imposition of stamp duty or the application of the exemptions, or to correctly calculate the stamp duty where it decided the land was dutiable. It exempted yellowacre on the basis that that it was primary production land. It concluded with the statement: 'Please note that these calculations are based on the provided Prolog rules and may not reflect actual legislation or any future changes. It is always recommended to consult with a legal professional for specific advice tailored to your situation.'

LLM	Description of output
BM06	The model failed to classify any of the land accurately for the imposition of stamp duty or the application of the exemptions, or to correctly calculate the stamp duty where it decided the land was dutiable. The explanations for the dutiability and the one application of the duty rates were confused and nonsensical. It concluded with the following statement: ‘If any questions about these analyses arise or if more specifics regarding each property would allow further detailed analysis on this subject matter are required we recommend consulting an official land valuer as they can provide definitive conclusions following specific Prolog rules application to a client’s unique case.’
CM07	The model truncated the question because it exceeded 2,000 characters. It was therefore inappropriate to proceed with the question.

Table 7: Stamp Duty Example – RAG with Prolog-Like Rules

LLM	Description of output
AM01	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemptions or calculate the duty where it contended it was payable. It stated that greenacre as residential land was exempt, greyacre was also exempt as residential land and yellowacre and brownacre were exempt as primary production land. It concluded: ‘The tax consequences of the conveyance of the identified land parcels vary based on their classification. It is important to consider the stamp duty implications before proceeding with the conveyance to ensure compliance with South Australian stamp duty regulations.’
AM02	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemptions. It exempted yellowacre and brownacre as primary production land. All duty computations were incorrect.
AM03	The model correctly classified the land other than brownacre as qualifying or non-qualifying land. It stated that brownacre was exempt as primary production land. It provided a stamp duty calculation for each piece of land it regarded as dutiable in a well laid out memorandum with a convenient summary of the tax consequences.
AM04	The model correctly classified the land other than blueacre as qualifying or non-qualifying land. It stated that blueacre was residential land. There were some duty computation errors. The output was well laid out and much simpler than AM03.
AM05	The model correctly classified the land other than blueacre as qualifying or non-qualifying land. The duty calculations were wrong and confused. The model’s conclusion then inconsistently stated that greyacre was exempt. It concluded: ‘Please note that this memorandum of advice is based on the

LLM	Description of output
	provided Prolog rules and may not reflect any additional or specific circumstances that may apply to each individual case.’
AM06	The model incorrectly classified the land other than blueacre as qualifying or non-qualifying land in its detailed analysis and some of the reasoning was confused. There were no duty calculations.
BM01	The model failed to classify any of the land accurately for the imposition of stamp duty or its exemption or calculate the duty where payable. It stated that greenacre and greyacre as residential land were exempt and yellowacre and brownacre as primary production were exempt. There was no attempt to calculate the duty.
BM02	The model failed to classify blackacre, blueacre, yellowacre or brownacre correctly. It stated that yellowacre and brownacre were exempt as primary production land. All duty computations were incorrect. It concluded: ‘Information is missing on the specific rates and provisions for duty rounding and possible concessions or exemptions beyond the classification-based exemptions provided in the contexts. Detailed provisions regarding specific conditions under South Australian law and possible additional considerations for mixed-use properties also were not provided.’
BM03	The model failed to classify blueacre as exempt and classified brownacre as exempt as primary production land. All duty computations were incorrect. The model concluded: ‘The stamp duty implications depend primarily on the classification of the land parcels in terms of residential and primary production uses. Taxes range from none to substantial based on these classifications and the property’s value.’
BM04	The model correctly classified the land other than blueacre as qualifying or non-qualifying land. It provided some incorrect stamp duty calculation for each piece of land that it regarded as dutiable.
BM05	The model correctly classified the land other than Blueacre as qualifying or non-qualifying land. It provided significantly incorrect stamp duty calculation for most pieces of land it regarded as dutiable.
BM06	The model failed to classify any of the land accurately for the imposition of stamp duty or the application of the exemption, or to correctly calculate the stamp duty. The explanations for the dutiability were confused and nonsensical. An example is the comment provided in respect of blackacre: ‘The stamp duty consequences of acquiring the Blackacre depends entirely on whether this piece is classified under residential, industrial use classifications only or some other form such that you cannot classify as being an agricultural one. Unfortunately without a clear determination what zone your parcel falls within it can be difficult to provide definitive advice here.’
CM07	As this model did not allow for a RAG file the question was not put to it.

APPENDIX 1 PART B**Stamp duty example: an output example – single shot prompt**

****Memorandum of Advice****

****Date: [Insert Date]****

****To: [Client's Name]****

****From: [Your Name]****

****Subject: Stamp Duty Consequences of Conveyance of Various Pieces of Land****

****Facts:****

- The following pieces of land are subject to conveyance:

1. Blackacre: Industrial land with a capital value of \$750,000.
2. Greenacre: Residential land with a capital value of \$2,000,000.
3. Blueacre: Commercial land with a capital value of \$3,000,000.
4. Greyacre: Residential land with a capital value of \$2,000,000.
5. Yellowacre: Primary production land with a capital value of \$4,000,000.
6. Brownacre: Land with no specific use, no improvements, and a capital value of \$350,000.

****Classification of Land:****

1. Blackacre: Classified as industrial land.
2. Greenacre: Classified as residential land with residential character.
3. Blueacre: Classified as commercial land with residential character.
4. Greyacre: Classified as residential land with minor improvements.
5. Yellowacre: Classified as primary production land with minor improvements.
6. Brownacre: Classified as land with no specific use and no improvements.

****Tax Consequences:****

- Stamp duty in South Australia is not imposed on qualifying land. The classification and use of the land play a significant role in determining the stamp duty consequences. Based on the provided information:

1. Blackacre (industrial land): Likely qualifies for stamp duty exemption.
2. Greenacre (residential land): Likely qualifies for stamp duty exemption.

3. Blueacre (commercial land): May not qualify for stamp duty exemption due to the commercial nature.

4. Greyacre (residential land): Likely qualifies for stamp duty exemption.

5. Yellowacre (primary production land): Likely qualifies for stamp duty exemption.

6. Brownacre (land with no specific use): Likely qualifies for stamp duty exemption.

****Computation of Tax Consequences:****

- For the pieces of land that qualify for stamp duty exemption, there would be no stamp duty payable upon conveyance. However, for land that does not qualify for exemption, the stamp duty amount would be computed based on the capital value of the land as per the relevant stamp duty rates in South Australia.

****Recommendation:****

- It is advisable to further verify the specific zoning and use classifications of each piece of land to determine the exact stamp duty consequences. Seeking professional advice from a conveyancer or tax advisor would be beneficial to ensure compliance with stamp duty regulations.

If you require any further assistance or clarification, please do not hesitate to contact me.

Yours sincerely,

[Your Name]

[Your Contact Information]

Stamp duty example: an output example – Prolog enhanced rules

Memorandum of Advice

****To:**** [Client Name]

****From:**** [Your Name/Your Firm]

****Date:**** [Insert Date]

****Subject:**** Stamp Duty Consequences of Land Acquisition in South Australia

****Facts****

You are considering the acquisition of the following parcels of land, each with distinct characteristics as described below:

1. ****Blackacre****: Industrial land with no improvements and no special zoning or determinations. Valued at \$750,000.

2. ****Greenacre****: Residential land with residential character improvements and a predominant use determination. Valued at \$2,000,000.

3. **Blueacre**: Commercial land with residential character zoning and no primary production determinations. Valued at \$3,000,000.
4. **Greyacre**: Residential land with minor improvements and residential zoning. Valued at \$2,000,000.
5. **Yellowacre**: Primary production land with minor improvements and residential zoning. Valued at \$4,000,000.
6. **Brownacre**: Unused land with no improvements, primary production zoning, and a primary production determination. Valued at \$350,000.

Analysis

The stamp duty consequences of acquiring each parcel of land depend on whether the land qualifies as “residential land,” “primary production land,” or other categories under South Australian law. Below, we analyze the classification and tax consequences for each parcel.

1. Blackacre

Classification:

- Use: Industrial.
- No residential or primary production characteristics.
- Not classified as residential land or primary production land.

Tax Consequences:

- Blackacre is **qualifying land** because it does not fall under the residential or primary production categories.
- **Stamp Duty:** Not payable.

2. Greenacre

Classification:

- Use: Residential.
- Improvements: Residential character.
- Determination: Predominant use for residential purposes.

Tax Consequences:

- Greenacre is **residential land** due to its residential use, improvements, and zoning.
- **Stamp Duty:** Payable.
- **Calculation:** Based on the value of \$2,000,000:
- Value exceeds \$500,000.

- Tax = $\$21,330 + (0.055 \times (\$2,000,000 - \$500,000))$

- Tax = $\$21,330 + \$82,500 = \$103,830$.

3. Blueacre

Classification:

- Use: Commercial.

- Zoning: Residential character zoning.

- No primary production determination.

Tax Consequences:

- Blueacre is **residential land** because it has residential character zoning.

- **Stamp Duty:** Payable.

- **Calculation:** Based on the value of \$3,000,000:

- Value exceeds \$500,000.

- Tax = $\$21,330 + (0.055 \times (\$3,000,000 - \$500,000))$

- Tax = $\$21,330 + \$137,500 = \$158,830$.

4. Greyacre

Classification:

- Use: Residential.

- Zoning: Residential.

Tax Consequences:

- Greyacre is **residential land** due to its residential use and zoning.

- **Stamp Duty:** Payable.

- **Calculation:** Based on the value of \$2,000,000:

- Value exceeds \$500,000.

- Tax = $\$21,330 + (0.055 \times (\$2,000,000 - \$500,000))$

- Tax = $\$21,330 + \$82,500 = \$103,830$.

5. Yellowacre

Classification:

- Use: Primary production.

- Zoning: Residential.
- No primary production determination.
- **Tax Consequences:****
- Yellowacre is ****primary production land**** due to its primary production use.
- ****Stamp Duty:**** Payable.
- ****Calculation:**** Based on the value of \$4,000,000:
 - Value exceeds \$500,000.
 - Tax = \$21,330 + (0.055 × (\$4,000,000 – \$500,000))
 - Tax = \$21,330 + \$192,500 = ****\$213,830****.

****6. Brownacre****

- **Classification:****
- Use: No specific use.
- Zoning: Primary production.
- Determination: Valuer General's primary production determination.
- **Tax Consequences:****
- Brownacre is ****primary production land**** due to its zoning and primary production determination.
- ****Stamp Duty:**** Payable.
- ****Calculation:**** Based on the value of \$350,000:
 - Value exceeds \$300,000 but not \$500,000.
 - Tax = \$11,330 + (0.05 × (\$350,000 – \$300,000))
 - Tax = \$11,330 + \$2,500 = ****\$13,830****.

****Summary of Tax Consequences****

<i>Land</i>	<i>Classification</i>	<i>Stamp Duty Payable</i>	<i>Amount (\$)</i>
-----	-----	-----	-----
<i>Blackacre</i>	<i>Qualifying Land</i>	<i>No</i>	<i>0</i>
<i>Greenacre</i>	<i>Residential Land</i>	<i>Yes</i>	<i>103,830</i>
<i>Blueacre</i>	<i>Residential Land</i>	<i>Yes</i>	<i>158,830</i>
<i>Greyacre</i>	<i>Residential Land</i>	<i>Yes</i>	<i>103,830</i>
<i>Yellowacre</i>	<i>Primary Production Land</i>	<i>Yes</i>	<i>213,830</i>

| *Brownacre* | *Primary Production Land* | *Yes* | 13,830 |

****Conclusion****

Based on the above analysis, stamp duty is payable on all parcels except Blackacre, which qualifies as exempt land. For the remaining parcels, the duty payable has been calculated in accordance with the applicable rates. Please let us know if further assistance is required.

Blueacre caused most of the model's difficulties and should be designated as exempt. The foregoing is easily corrected to address that error.

APPENDIX 1 PART C: STAMP DUTY EXAMPLE – PROLOG-LIKE RULES

%Residential land

%land(,,,_).

%residential land

residential_use(Identifier,Use,Improvements,Zoning,):-

land(Identifier,residential,_,_);

land(Identifier,_,residential_character_improvements,_,_);

land(Identifier,_,_,solely_residential_zoning,_,_).

%primary production land

primary_production_use(Identifier,Use,_,_,Valuer_General_Primary_Production_Determination):-

land(Identifier,primary_production,_,_);

land(Identifier,_,_,primary_production).

% dutiable – it is not qualifying land

stamp_duty_payable(Identifier,_,_,_):-

land(Identifier,_,_,_),

residential_use(Identifier,_,_,_);

primary_production_use(Identifier,_,_,_).

% not dutiable – it is qualifying land

stamp_duty_not_payable(Identifier,_,_,_):-

land(Identifier,_,_,_),

\+ residential_use(Identifier,_,_,_),

```

\+ primary_production_use(Identifier,_,_,_).

#####

% Stamp Duty Rates
% The legislation requires the amount to be rounded up the amount to the next $100 – this
does not do, it only calculates the rates
%Case1: Amount does not exceed $12,000
stampduty_rate_tax(Amount, Tax) :-
    Amount <= 12000,
    Tax is 0.01 * Amount.

% Case 2: Amount exceeds $12,000 but not $30000
stampduty_rate_tax(Amount, Tax):-
    Amount > 12000,
    Amount <= 30000,
    Excess is Amount - 12000,
    Tax is 120 + (0.02 * Excess).

% Case 3: Amount exceeds $30,000 but not $50,000
stampduty_rate_tax(Amount, Tax):-
    Amount > 30000,
    Amount <= 50000,
    Excess is Amount - 30000,
    Tax is 480 + 0.03 * Excess.

% Case 4: Amount exceeds $50,000 but not $100,000
stampduty_rate_tax(Amount, Tax):-
    Amount > 50000,
    Amount <= 100000,
    Excess is Amount - 50000,
    Tax is 1080 + 0.035 * Excess.

% Case 5: Amount exceeds $100,000 but not $200,000
stampduty_rate_tax(Amount, Tax):-

```

Amount > 100000,

Amount =< 200000,

Excess is Amount - 100000,

Tax is $2830 + 0.04 * \text{Excess}$.

% Case 6: Amount exceeds \$200,000 but not \$250,000

stampduty_rate_tax(Amount, Tax):-

Amount > 200000,

Amount =< 250000,

Excess is Amount - 200000,

Tax is $6830 + 0.0425 * \text{Excess}$.

% Case 7: Amount exceeds \$250,000 but not \$300,000

stampduty_rate_tax(Amount, Tax):-

Amount > 250000,

Amount =< 300000,

Excess is Amount - 250000,

Tax is $8955 + 0.0475 * \text{Excess}$.

% Case 8: Amount exceeds \$300,000 but not \$500,000

stampduty_rate_tax(Amount, Tax):-

Amount > 300000,

Amount =< 500000,

Excess is Amount - 300000,

Tax is $11330 + 0.05 * \text{Excess}$.

% Case 8: Amount exceeds \$500,000

stampduty_rate_tax(Amount, Tax):-

Amount > 500000,

Excess is Amount - 500000,

Tax is $21330 + 0.055 * \text{Excess}$.

APPENDIX 2 PART A**Land tax example: an output example – Prolog enhanced rules****MEMORANDUM OF ADVICE - LAND TAX IMPLICATIONS****1. FACTS****Current Land Holdings:**

- Jim personally owns Blackacre (commercial) - \$2,000,000
- Mary personally owns Whiteacre (commercial) - \$1,500,000
- Jim's Family Trust owns Brownacre (commercial) - \$3,000,000
- Mary's Trust owns:
 - * Greenacre (primary production) - \$2,500,000
 - * Yellowacre (residential) - \$650,000
- Jim Mary Trust owns Cityblock (commercial) - \$4,500,000
- May owns Blueacre (principal residence) - \$1,250,000

Proposed Acquisition:

- Highville (commercial shopping centre) - \$3,250,000

2. CURRENT TAX CONSEQUENCES**Individual Holdings:**

- Jim: \$2,000,000 - Land tax = \$28,110 + $(2.4\% \times \$0) = \$28,110$
- Mary: \$1,500,000 - Land tax = \$7,570 + $(2.0\% \times \$0) = \$7,570$

Trust Holdings:

- Jim's Family Trust: \$3,000,000 - Land tax = \$40,773 + $(2.4\% \times \$262,000) = \$47,061$
- Mary's Trust: \$3,150,000 - Land tax = \$40,773 + $(2.4\% \times \$412,000) = \$50,661$
- Jim Mary Trust: \$4,500,000 - Land tax = \$40,773 + $(2.4\% \times \$1,762,000) = \$83,061$

3. TAX CONSEQUENCES OF ACQUIRING NEW PROPERTY

Adding Highville (\$3,250,000) to existing entities would result in:

- Jim: Total \$5,250,000 - Significant increase in land tax
- Mary: Total \$4,750,000 - Significant increase in land tax
- Existing Trusts: Would substantially increase their already high tax burden

4. RECOMMENDATION

I recommend establishing a new trust to acquire Highville. Reasons:

1. A new trust would be assessed independently, avoiding aggregation with existing properties
2. The trust would pay land tax on \$3,250,000 which would result in:

$$\text{Tax} = \$40,773 + (2.4\% \times \$512,000) = \$53,061$$

This represents the lowest possible land tax outcome compared to adding the property to any existing entity. The new trust should be structured to ensure it's not grouped with existing trusts for land tax purposes, while still maintaining the desired family control and beneficial interests.

The new structure would result in the lowest ongoing land tax cost while maintaining family control over the asset.

(The foregoing analysis has an error. It includes Greenacre in the computations when it is primary production land and exempt. Once again that is reasonably easily corrected. As also discussed earlier, a company or one of the children would achieve an ongoing lower land tax consequence. However, for other tax reasons a new discretionary trust is more than likely the preferred entity.)

APPENDIX 2 PART B: LAND TAX EXAMPLE – PROLOG-LIKE RULES

% South Australian land tax rates

% Standard Rates - These rates apply to companies and individuals

%Base case: Amount does not exceed \$732,000

standard_rate_tax(Amount, 0) :-

Amount =< 732000.

% Case 1: Amount exceeds \$732,000 but not \$1,176,000

standard_rate_tax(Amount, Tax) :-

Amount > 732000,

Amount =< 1176000,

Excess is Amount - 732000,

Tax is 0.5 * ceiling(Excess / 100).

% Case 2: Amount exceeds \$1,176,000 but not \$1,711,000

```

standard_rate_tax(Amount, Tax) :-
    Amount > 1176000,
    Amount <= 1711000,
    Excess is Amount - 1176000,
    Tax is 2220 + (1.0 * ceiling(Excess / 100)).

% Case 3: Amount exceeds $1,711,000 but not $2,738,000
standard_rate_tax(Amount, Tax) :-
    Amount > 1711000,
    Amount <= 2738000,
    Excess is Amount - 1711000,
    Tax is 7570 + (2.0 * ceiling(Excess / 100)).

% Case 4: Amount exceeds $2,738,000
standard_rate_tax(Amount, Tax) :-
    Amount > 2738000,
    Excess is Amount - 2738000,
    Tax is 28110 + (2.4 * ceiling(Excess / 100)).

% Trust cases - These rates apply to trusts
% Base case: Amount does not exceed $25,000
trust_tax(Amount, 0) :-
    Amount <= 25000.

% Case 1: Amount exceeds $25,000 but not $732,000
trust_tax(Amount, Tax) :-
    Amount > 25000,
    Amount <= 732000,
    Excess is Amount - 25000,
    Tax is 125 + (0.5 * ceiling(Excess / 100)).

% Case 2: Amount exceeds $732,000 but not $1,176,000
trust_tax(Amount, Tax) :-

```

Amount > 732000,

Amount =< 1176000,

Excess is Amount - 732000,

Tax is 3660 + (1.0 * ceiling(Excess / 100)).

% Case 3: Amount exceeds \$1,176,000 but not \$1,711,000

trust_tax(Amount, Tax) :-

Amount > 1176000,

Amount =< 1711000,

Excess is Amount - 1176000,

Tax is 8100 + (1.5 * ceiling(Excess / 100)).

% Case 4: Amount exceeds \$1,711,000 but not \$2,738,000

trust_tax(Amount, Tax) :-

Amount > 1711000,

Amount =< 2738000,

Excess is Amount - 1711000,

Tax is 16125 + (2.4 * ceiling(Excess / 100)).

% Case 5: Amount exceeds \$2,738,000

trust_tax(Amount, Tax) :-

Amount > 2738000,

Excess is Amount - 2738000,

Tax is 40773 + (2.4 * ceiling(Excess / 100)).

APPENDIX 3 PART A**CGT Division 152 example: an output example – Prolog enhanced rules**

MEMORANDUM OF ADVICE

Date: [Current Date]

Re: Jim's Eligibility for Division 152 CGT Relief

FACTS:

- Jim has sold his business
- CGT Event A1 has occurred
- Capital gain of \$1,900,000
- Income history:
 - * Two years prior: \$1,200,000
 - * Last year: \$1,400,000
 - * Current year: \$1,000,000
- Gross assets: \$8,000,000
- The asset is confirmed to be an active asset
- No partnership involvement
- Conditions 1A/1B not applicable
- Not a CGT asset or asset partnership

ANALYSIS OF BASIC CONDITIONS:

1. CGT Event Test (Test A):

✓ SATISFIED

- A CGT event A1 has occurred (business sale)

2. Capital Gain Test (Test B):

✓ SATISFIED

- The gain of \$1,900,000 is greater than zero

3. Small Business Test (Test C):

This can be satisfied through either:

a) Small Business Income Test:

✓ SATISFIED

- All years' income is below \$2,000,000 threshold:

* Two years prior: \$1,200,000 < \$2,000,000

* Last year: \$1,400,000 < \$2,000,000

* Current year: \$1,000,000 < \$2,000,000

b) Gross Assets Test:

✗ NOT SATISFIED

- Gross assets of \$8,000,000 exceed the \$6,000,000 threshold

4. Active Asset Test (Test D):

✓ SATISFIED

- The asset is confirmed as an active asset

BASIS FOR DECISION:

According to the Prolog rules provided, Jim needs to satisfy all four basic conditions (Tests A, B, C, and D) to qualify for Division 152 relief. While he satisfies Tests A, B, and D, Test C can be satisfied through multiple pathways, and Jim satisfies this through the small business income test, despite not meeting the gross assets test.

The capital gain relief rule requires all conditions to be met, which they are in Jim's case through:

- Valid CGT event (A1)

- Positive capital gain

- Meeting the small business test through income requirements

- Active asset confirmation

CONCLUSION:

Based on the analysis above, Jim DOES satisfy the basic conditions for Division 152 relief. He is therefore entitled to the additional 50% capital gains discount, subject to any other specific requirements not covered in this basic conditions analysis.

Note: This advice is limited to the basic conditions analysis only. Jim should seek additional advice regarding any specific requirements or obligations related to claiming the Division 152 relief.

This model did not directly address the partnership asset alternative or subsections 1A and 1B in Test C, but it interpreted the facts correctly and that these requirements were irrelevant.

APPENDIX 3 PART B

CGT Division 152 example – Prolog-like rules

% As the question focuses on a non partnership interest the partnership aspects and subsections 1A or 1B requirements have not been independently tested.

capitalgainrelief(Entity, CGTEvent, Gain, TwoYearsPriorIncome, LastYearIncome, CurrentYearIncome, GrossAssets, Partner, Conditions1A1B, CGTAsset, AssetPartnership, ActiveAsset):-

div152_facts(Entity,_,_,_,_,_,_,_,_,_,_),

```
cgt_event_test_a(Entity,_,_,_,_,_,_,_,_,_,_,_),
```

cgt_gain_test_b(Entity,_,_,_,_,_,_,_,_,_,_),

```
cgt_gain_test_c(Entity,_,_,_,_,_,_,_,_,_,_),
```

```
cgt_activeasset_test_d(Entity,_,_,_,_,_,_,_,_,_,_),!.
```

```
cgt_event_test_a(Entity,_,_,_,_,_,_,_,_,_,_):-
```

div152_facts(Entity, a1,_,_,_,_,_,_,_,_,_).

```
cgt_gain_test_b(Entity,_,Gain,_,_,_,_,_,_,_):-
```

div152 facts(Entity, ,Gain, , , , , , , ,),

Gain>0.

```
cgt_gain_test_c(Entity,_,_,_,_,_,_,_,_,_,_):-
```

```
cgt_small_business_test(Entity, , , , , , , , , );
```

```
cgt_gross_assets_test(Entity,_,_,_,_,_,_,_,_,_);
```

```
cgt_partner_test(Entity,_,_,_,_,_,_,_,_,_,_);
```

cgt condition1A1B test(Entity, , , , , , , , , , ,).

cgt small business test(Entity, , , , , , , , , , ,):-

cgt twoprior_year income test(Entity,, , , , , , , , ,),

cgt_last_year income test(Entity, , , , , , , , , , ,),

cgt current year income test(Entity, , , , , , , , , ,);

cgt current year income test(Entity, , , , , , , , , , ,).

```
cgt twoprior year income test(Entity, , TwoYearsPriorIncome, , , , , , ):-
```

div152 facts(Entity, , ,TwoYearsPriorIncome, , , , , , , ,),

TwoYearsPriorIncome < 2000001.

cgt_last_year_income_test(Entity,_,_,_,LastYearIncome,_,_,_,_,_):-

div152_facts(Entity,_,_,_,LastYearIncome,_,_,_,_,_),

LastYearIncome < 2000001.

cgt_current_year_income_test(Entity,_,_,_,CurrentYearIncome,_,_,_,_,_):-

div152_facts(Entity,_,_,_,CurrentYearIncome,_,_,_,_,_),

CurrentYearIncome < 2000001.

cgt_gross_assets_test(Entity,_,_,_,GrossAssets,_,_,_,_,_):-

div152_facts(Entity,_,_,_,_,GrossAssets,_,_,_,_,_),

GrossAssets < 6000001.

cgt_partner_test(Entity,_,_,_,_,Partner,_,CGTAsset,AssetPartnership,_)ate:-

div152_facts(Entity,_,_,_,_,_,yes,_,_,_,_),

div152_facts(Entity,_,_,_,_,_,_,yes,_,_,_),

div152_facts(Entity,_,_,_,_,_,_,_,yes,_,_,_).

cgt_condition1A1B_test(Entity,_,_,_,_,_,Conditions1A1B,_,_)ate:-

div152_facts(Entity,_,_,_,_,_,_,yes,_,_,_,_).

cgt_activeasset_test_d(Entity,_,_,_,_,_,_,_,ActiveAsset)ate:-

div152_facts(Entity,_,_,_,_,_,_,_,_,yes).