Building Sustainable Free Legal Advisory Systems: Experiences from the History of AI & Law

GRAHAM GREENLEAF, ANDREW MOWBRAY AND PHILIP CHUNG

(2018) 34(1) *Computer Law & Security Review*
[2017] *UNSWLRS* 53

UNSW Law
UNSW Sydney NSW 2052 Australia

Building sustainable free legal advisory systems: Experiences from the history of AI & law

Graham Greenleaf, Andrew Mowbray and Philip Chung[*]

14 February 2018 – Acknowledgments[**]

Accepted for publication in (2018) 34(1) *Computer Law & Security Review*

Abstract

The enthusiasm for artificial intelligence (AI) as a source of solutions to problems is not new. In law, from the early 1980s until at least the early 2000s, considerable work was done on developing 'legal expert systems.' As the DataLex project, we participated in those developments, through research and publications, commercial and non-commercial systems, and teaching students application development. This paper commences with a brief account of that work to situate our perspective.

The main aim of this paper is an assessment of what might be of value from the experience of the DataLex project to contemporary use of 'AI and law' by free legal advice services, who must necessarily work within funding and other constraints in developing and sustaining such systems. We draw fifteen conclusions from this experience which we consider are relevant to development of systems for free legal advice services. The desired result, we argue, is the development of integrated legal decision-support systems, not 'expert systems' or 'robot lawyers'. We compare our insights with the approach of the leading recent text in the field, and with a critical review of the field over twenty-five years. We conclude that the approach taken by the DataLex project, and now applied to free legal advice services, remains consistent with leading work in field of AI and law.

The paper concludes with brief suggestions of what are the most desirable improvements to tools and platforms to enable development of free legal advice systems. The objectives of free access to legal information services have much in common with those of free legal advice services. The information resources that free access to law providers (including LIIs) can provide will often be those that free legal advice services will need to use to develop and sustain free legal advisory systems. There is therefore strong potential for valuable collaborations between these two types of services providers.

[*] Respectively, Professor of Law & Information Systems, UNSW Australia and Co-founder, AustLII; Professor of Law & Information Technology, University of Technology, Sydney (UTS) and Co-Director, AustLII; and Associate Professor of Law, UNSW Australia and Executive Director, AustLII.

[**] Thanks to Roger Clarke, Anna Cody, George Williams, Lyria Bennett Moses, Matthew Keiley, Jill Matthews and three referees for valuable comments; all content remains the responsibility of the authors.

Contents

1 Changing perspectives on AI and law

1.1 Waves of enthusiasm

The enthusiasm for artificial intelligence (AI) as a source of solutions to problems in all walks of life, including law, is not new. The first wave of excitement about AI-based solutions to a general range of problems, from the early 1980s, focussed primarily on 'legal expert systems' – that is, systems capable of rendering advice concerning specific types of problems in narrowly defined legal domains. It included a previous iteration of computing practices that were going to 'make everything different' such as Japan's 'Fifth Generation' project. Many believed this project would make Japan dominant in computing within a decade, because parallel processors and logic programming techniques (in which Japan was believed to have a lead) would enable inferences to be drawn from (newly possible) 'massive' databases of domain-specific knowledge. The revolution failed to launch,[1] but more modest versions of the expert system boom continued until the late 90s,[2] by which time enthusiasm for the Internet attracted people's interest and greed (the first '.com' boom) in substitution. Many valuable expert systems, academic, commercial and governmental, were developed along the way.[3] Attempts have been made to catalogue the most valuable 'AI and law' research of the period.[4]

In this paper we do not attempt to add to the existing reviews of the field of AI and law in its earlier wave, nor attempt to determine whether the new enthusiasms are fundamentally different. We primarily wish to reflect on our own participation in it as 'the DataLex Project', through which we were active participants both academically (research and teaching) and commercially, from 1984 to 2001. In this paper we attempt to distil nearly two decades of experience, because some of what we learned may still be of value. Not all experience turns out to be valuable, but proceeding in ignorance of it can often be a waste of time and money. We compare our insights with the approach of the leading recent text by a proponent of this field of research (Ashley[5]), and with a critical review of the whole field over twenty-five years (Leith). [6]

We also wish to focus on the particular issues facing one particular group of users of 'AI and law' technologies, namely services which provide free legal advice (as discussed below): we do not focus on the issues that are of more relevance to financially well-endowed commercial organisations. In the concluding parts of the paper we consider what desirable improvements to development of free legal advisory systems might now be more feasible than before, and how these could lead to new collaborations between 'free access to law' providers (such as legal information institutes) and providers of free legal advice.

1.2 A different focus: 'Sustainable free legal advisory systems'

Organisations that provide free legal advice services, such as community legal centres and advice bureaus, often have large numbers of part time volunteers and interns providing legal advice and assistance, sometimes relatively early in their legal careers. Yet they are required to provide legal advice on a very wide range of legal issues, often with high client case-loads, and sometimes with numerous instances of similar problems (for example, consumer law or

[1] See Feigenbaum and McCorduck, 1983; Pollack, 1992. For an overview, see Wikipedia, 'Fifth Generation Computer' (29 July 2014).
[2] DataLex, 1989c discusses the various 'waves' of 'AI and law'.
[3] One of the earliest and best known was Capper and Susskind, 1988; SoftLaw Corporation developed decision-support systems in the early 1990s on a much larger scale, used for large-scale decision-making in Australia's social security and veterans' affairs administrations.,
[4] Bench-Capon et al, 2012
[5] Ashley, 2017; Professor Kevin Ashley has been a major contributor to the field since the 1980s and is co-Editor-in-Chief of its leading journal, *Artificial Intelligence and Law*.
[6] Leith, 2010; Professor Philip Leith was a significant contributor to the field from its earliest stages, including in the utilisation of Prolog for legal representation, and is the author of many articles in the field.

immigration law issues). They usually have few financial resources to be able to utilise outside counsel (although pro bono schemes assist in this), or simply to hire more staff. Commercial legal publishers often do not publish in the areas of law on which these services need expert information, particularly for marginalised clienteles. So they must develop their own resources, sometimes in the form of plain English 'law handbooks'.

Knowledge-based advisory systems may be able to assist these often 'semi-expert' legal advisors to meet the heavy demands placed on them. The constraints on government legal aid services are different, but these systems may also be of value to them. We must distinguish between systems which are to be used only by those who work for such services (whether as employees, interns or pro-bono), from those systems which are made available to their clients for their direct use (client end-users). Such systems are likely to be significantly different. Self-represented litigants, and members of the public attempting to resolve their own legal issues without resorting to litigation, are another 'end-user client' category where knowledge-based advisory systems might play a valuable role. In this paper, we refer to all these types of services as 'free legal advice services', and knowledge-based systems to assist them, whether used only by the services or used by end-user clients, as 'free legal advisory systems.'

There are close relationships between free legal advice services and providers of free access to legal information (LIIs), such as AustLII,[7] which have as a main motivation the provision of free access to the general public, not only to the legal profession or universities. LIIs see community legal services as intermediaries which provide the benefits of free access legal information to the public, and thus one of the main audiences for the information they provide. Such services, often of necessity, are substantial users of LIIs and other sources of free access legal information. There are also close ties between the free legal information services and technologies provided by organisations like AustLII and other LIIs, and the information and technologies needed to created effective legal advisory systems.

Both LIIs and free legal advice services often face similar constraints: they need to operate at a relatively low cost; they have few funds to purchase outside services; they rely to a significant extent on open source software; and maintenance costs are a significant issue. Free legal advice services often have only modest levels of technical expertise on which to draw.

The result, as we see it, is that if free legal advice services are to make use of knowledge-based technologies, they need to find software and knowledge-based applications that are not only low cost (preferably free, possibly open source), but such software and applications must be maintainable from their internal resources, because external maintenance will be prohibitively expensive. We describe this goal as 'sustainable legal advisory systems'. Assisting their development and use fits the mission of a LII like AustLII,[8] so it is also appropriate for a LII to be involved in free legal advisory systems. The purpose of this paper is to explore what our (and AustLII's) experience in this field can bring to the achievement of this goal.

2 Personal experiences: The DataLex Project

From its earliest years we were active participants in the AI and law field, both as developers of the original DataLex Project (1984-2001) and as the authors of scholarship on legal expert systems. Like the Australasian Legal Information Institute (AustLII) which grew out of it, it was a joint project of academics from UNSW Australia and from the University of Technology, Sydney. A more detailed history, and full bibliography, of the DataLex project is available

[7] Australasian Legal Information Institute <http://www.austlii.edu.au>
[8] See *AustLII Annual Report* 2016, pp. 6-8, for a statement of AustLII's mission and objectives.

elsewhere,[9] but a brief account is necessary here, because that history is the basis of the arguments we put forward in the next section on the uses of AI in law. The Project went through four stages:

Pre-web, pre-AustLII, pre-Workstations (1984-90) DataLex began as a project focusing solely on legal expert systems, with development of the LES shell for procedural (decision network) inferencing.[10] Our initial ambition was to assist community legal centres,[11] but events led the project down other research, legal education and commercial paths. A case-based inferencing mechanism based on nearest-neighbour discriminant analysis (PANNDA), developed by Alan Tyree, was added[12]. Further software development then added a full text retrieval system, AIRS, which emulated the STATUS retrieval system,[13] and a hypertext engine (HYPE) which was an early pre-web development of this genre of software.[14] Initial applications were expert systems on intestacy law (INTEST[15]) and copyright law (COPYRITA[16]). FINDER answered questions about ownership of found objects based on case law.[17] A legal information retrieval training system used AIRS to simulate the performance of the CLIRS commercial system, with small databases[18] and was purchased by Australian law schools. A textbook supported it.[19] A commissioned hypertext demonstration using HYPE, plus AIRS text retrieval, ran over a remote dial-up (LAWS OF AUSTRALIA Defamation Demonstration[20]).

DataLex Workstations (1990-95) In the years shortly before the development of the World-Wide-Web (1990-93), the DataLex approach was based on the integration of inferencing (primarily rule-based and to some extent case-based expert systems), hypertext and text retrieval, with some document generation capacity as well. This first required development of rule-based inferencing software (XSH) (backward and forward chaining rule-based reasoning), subsequently refined by the addition of quasi-natural-language knowledge representation (YSH), influenced by the SoftLaw systems.[21] . The key pieces of software (XSH and later YSH, HYPE and AIRS),[22] were integrated into the 'DataLex Workstations' approach.[23] From 1990 a commercial avatar of the project, DataLex Pty Ltd, developed and licensed 'workstations', primarily the 'Intellectual Property Workstation' and the 'Privacy Workstation',[24] which used all three technologies. Updates to the content were distributed on disk, in the absence of any effective online alternative. The Workstations had modest commercial success. Support stopped in 1995, when a commercial publisher terminated DataLex's licence to include case law content. The software and development approach was taught at UNSW and UTS. DataLex Pty Ltd also carried out consultancy work on re-development of commercial and government systems, but AustLII's arrival soon made these systems increasingly redundant.

---

[9] G Greenleaf, A Mowbray and P Chung The DataLex Project: History and Bibliography, 2017 <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3095897>

[10] DataLex, 1995, 1985b; The STATUS search software was used on the CLIRS system, subsequently Info-One, from 1985.

[11] Assisting community legal services is where the DataLex Project intended to start in 1985, with a funding application to the NSW Law Foundation to develop a 'database of legal expert systems' to assist such services, based on the approach that 'expertise is relative'.

[12] DataLex, 1985a - PANNDA (precedent analysis by nearest-neighbour discriminant analysis) inferencing system

[13] DataLex, 1987b

[14] DataLex, 1989a

[15] DataLex, 1985b

[16] DataLex, 1987c

[17] DataLex, 1985c, 1986,1989d

[18] DataLex, 1988c

[19] DataLex, 1988b

[20] DataLex, 1989b

[21] Mead and Johnston, 1991

[22] Those programs, and their integration, were authored by Andrew Mowbray.

[23] DataLex, 1992a, 1992b, 1995

[24] DataLex, 1991a, 1992

The AustLII context – Web 1.0 (1996-2001) The DataLex project, and the 'Workstations' developed under it, had a very substantial influence on the techniques and approach implemented in the development of AustLII from 1995,[25] particularly the development of HYPE into a tool for automated generation of large-scale automated hypertext mark-up of legal documents. From 1995, development of AustLII required concentration on text retrieval and hypertext, and further development of those aspects stemming from the DataLex Project became part of AustLII's ongoing development.[26] For the first six years of AustLII (1995-2001) the inferencing aspects of the DataLex Project were transferred to the new web environment, the integration of inferencing (knowledge-bases and dialogues) with hypertext and text retrieval was further developed, and methods of 'collaborative inferencing' (distributed, multi-author knowledge-bases) were pioneered, all in the AustLII context.[27] This may have been the first significant attempt to develop legal inferencing systems on the web.[28]

From 2001 (and the Web 2.0 era), AustLII did not focus on inferencing and knowledge-bases, and its 'AI-related' work instead concentrated on the use of heuristics to improve text retrieval and hypertext mark-up, and also the use of heuristics for automated construction of an international case and journal citator (LawCite).[29] AI techniques have therefore continued to be essential to AustLII's work. All of the DataLex Project software continues to operate, and is again being actively developed, as interest has emerged in the new wave of 'AI and law'.

3 Experience relevant to free legal advice systems

The current wave of enthusiasm for 'AI and Law' takes many different forms: machine learning from training sets in areas such as document discovery; prediction of litigation outcomes; 'smart contracts'; and advisory systems – to name but a few. Greater computing capacities, new technologies and better interfaces all contribute to create new opportunities, but their very variety creates a serious need for the ability to discriminate between what is on offer. To do so effectively it is important to realise that this field is not a *tabula rasa*.

The original DataLex work was created prior to broad access to the World-Wide-Web (1993 onward), and prior to free access to law. It was therefore prior to any feasible ideas of inferencing systems being available for free or as part of free access online legal services. Even very low cost distribution by dial-up systems, CD-ROM or floppies was not very feasible. Nevertheless, most of the conclusions we reached in articles published from 1987-2001 continue to have validity in relationship to Internet-distributed systems. There are numerous other authors, both at that time and since, who have advocated versions of the 'decision support system' approach that we outline below, as discussed by Paliwala, 2016.

The rest of this section documents the main conclusions we reached from our experience in developing and using the DataLex software and applications, with references to our publications. We have avoided importing references to post-2001 (ie Web 2.0) technologies or publications, so as to better reflect this experience. Although this is a lengthy fifteen item list, it is summarised in [3.16].

3.1 Law is not 'just another problem domain'

Non-lawyers too easily assume that law is much the same as any other subject domain in which AI tools can be used. Worse, they assume that law is an easy domain, because 'its

[25] DataLex, 1995, 1995a
[26] DataLex, 1995a
[27] DataLex, 1997, 1997a, 1997b, 2000
[28] DataLex, 1997 – part 4
[29] Mowbray, Chung and Greenleaf, 2016

provisions are written down'.[30] To the contrary, law is an unusual and difficult problem domain, because legal expert systems do not usually involve modelling either (i) heuristics of how experts make decisions, or (ii) causal models of physical systems. The relationship between legal sources and legal reasoning is unique. The characteristics of law that make it a particularly challenging domain for development of expert systems were well recognised by some authors from the mid-1980s,[31] and ignored by others.[32]

The types of legal expertise that may need to be modelled include: general domain knowledge; formal knowledge of legal sources; logical reasoning; interpretative skills; research skills; organisational skills; strategic skills; communication skills; and 'real world' knowledge. Depending on what types of expertise a system is intending to capture, at least five main types of systems may be built: 'formal' advisory systems; 'strategic' advisory systems; automatic document generators; 'intelligent' litigation / transaction support systems; and 'intelligent' retrieval systems.[33] 'Formal' advisory systems are characterised by the necessity for any conclusions reached by the system to be supported by legally convincing reasons (and 'correctness' of the advice has little meaning in the legal context beyond 'convincing'). The rest of these suggestions in [3.2]-[3.15] are primarily relevant to such formal advisory systems.

What we call 'strategic' advisory systems, Ashley describes as the 'expert systems' paradigm (which he rejects): developed from manual questioning of human experts concerning problem scenarios, and capturing the rules they say they use to solve them.[34] Such expert systems based on expert heuristics (rules of thumb) are certainly one classic form of expert system, but 'legal expert systems' always included what we call 'formal' advisory systems, based primarily on representing formal legal doctrine in rules (according to expert interpretations), with lawyer's heuristics playing only a minor role. Leith, while using different terminology, seems to agree that what we call 'strategic' and 'formal' advisory systems are the 'two basic … options' of 1980s legal expert systems development (both of which he rejected).[35] As we argue later, we agree that both these options are of little use by themselves.

3.2 Expertise is relative

What counts as a legal 'expert' system depends on the context. If we assume a lack of expertise in bankruptcy law in volunteers in a community legal centre, a system might give valuable (and otherwise unaffordable) advice on bankruptcy law in that context, even if it was limited in its capacity to deal with complex bankruptcy issues. The same system would not be regarded as 'expert' (or useful) in the context of a law firm specialising in finance law.[36] Any well-designed system will attempt to identify those situations where the best advice it can give is 'this problem appears to beyond the expertise of this system' and refers the user to the appropriate human expert. Where legal advice services have large numbers of clients seeking advice, systems with such limited expertise could (in the right situations) be valuable as a method to (in effect) triage which clients need to see a lawyer, or have the most urgent need.

The rest of the argument of this article primarily concerns factors that are relevant to the building of systems for organisations equivalent to community legal centres, assuming they have quite limited funds to either build or (more importantly) maintain AI-based legal resources: [1.2]. In particular, they can be assumed to have only very limited access to

[30] DataLex, 1987a, pp. 2-6, 'Law and Expert Systems'; DataLex 1989c, 'Limitations on automated legal reasoning'
[31] Waterman et al, 1986; Susskind, 1986; DataLex, 1987a, 1987b, 1989c, 1989d
[32] Michie and Johnson, 1984; Sergot and Kowalski, 1986
[33] DataLex, 1989c, 'Types of legal expertise' and 'Types of knowledge-based applications to legal practice'
[34] Ashley, 2017, p. 8.
[35] Leith, 2010, p. 4.
[36] DataLex, 1987a, p. 8, '4 "Basic DataLex Design Criteria: Expert" is a relative notion'

computing experts with experience in AI tools or AI development. This practical limitation tends to be ignored in theoretical discussions of systems development, where the assumption of unlimited development resources is easily made, but obscures the reality of the range of developments that may be of utility.

3.3 No AI tool suits all types of problems

Two fundamental questions are: is the system which is being built intended to provide justifications for its answer based on underlying legal sources (legislation, cases etc): or is it intended to give the 'correct' provision of useful answers without legal reasons (eg should a document be discovered; or what is the percentage likelihood of success of a particular piece of litigation).[37] If the answer is the former, the tools which can be used are limited to those which include and use knowledge representations of those sources. As discussed in [4.1], other systems may attempt to model other forms of legal expertise. The rest of this discussion assumes (as was always the primary focus of the DataLex Project[38]) that the intention is to build systems which can justify their answers/recommendations by reference to the relevant law.

For example, some AI tools are based on machine learning, but for many legal problems, there are no sufficiently comprehensive 'training sets' of previous examples which have dealt with the same problem. For document discovery systems, training sets are easy to construct, and the costs of doing so easy to justify. Ashley explains that IBM's Watson, a current exemplar of AI, cannot (at least yet) answer 'legal questions' because 'one expects an explanation of why the answer is well-founded', not just a correct answer.[39] However, systems like Watson (or its legal derivative, Ross), can be trained to recognise different forms of the same question,[40] which can be valuable in new forms of decision support systems.

3.4 Start with legislation, not case law

Despite nearly 40 years of research into case-based legal reasoning, there is probably not much role for AI representations of case law as yet, at least in the types of systems that might be applied and sustainable by community legal advice services, and which aim to provide conclusions by reference to legal principles. The theoretical basis of case-based reasoning is still too much in dispute, and the identification of attributes in reported cases is still too time consuming and requires too much expertise to represent (and therefore is too costly). In some narrow legal domains with high financial returns, the position may be different. Although the DataLex Project experimented with case-based reasoning, only a very limited range of applications resulted.[41]

'Ready-made' or easily captured large-scale 'training sets' of case law are generally not available, so some AI tools which are very successful in some contexts (eg document discovery systems), do not provide general solutions in other problem areas. Predicting outcomes of litigation using correlations with factors that have no direct relationship to formal reasoning is increasingly possible using 'big data' analytics (what used to be called 'jurimetrics'), but is not likely to have any sustainable use by free legal advice services.

A more feasible approach toward use of case law is to attempt to alert users of inferencing systems when case law reasoning is required, and assist them to get to the right starting points to find the most relevant cases, including by access to expertise captured in

---

[37] DataLex, 1987a, p. 5, 'Deep v Shallow Models'
[38] DataLex, 1987a, p. 7, 'Basic DataLex Design Criteria: 3 Full Justification of Advice Must be Provided'
[39] Ashley, 2017, p. 17.
[40] Ashley, 2017, p. 152, concerning Ross.
[41] DataLex, 1988a; 1985c; 1989d; 1995: '6.2 Examples and the inference engine'

commentary: see [3.11] – [3.13]. While Ashley also considers that this is an important objective,[42] he and other researchers in case-based reasoning also have more ambitious goals.

In our view, the best starting points for AI-based legal systems, at least for free access legal advice providers of limited means, are still the representation of legislation, or procedural problems which depend to a large extent on legislation. For other use cases, more ambitious (and expensive) approaches to AI representations of case law may be justifiable.

One of the limitations of Ashley's 'cognitive computing paradigm' is that 'the knowledge is embodied in the corpus of texts from which the program extracts candidate solutions or solution elements and ranks them in terms of their relevance to the problem,'[43] by using machine learning (ML) to automatically extract and rank candidate solutions from the corpus. He assumes that 'if the problem is a fact situation about which to make arguments concerning a legal claim, a corpus of legal cases involving that type of claim would be required'.[44] The shortcoming of this argument is that a very significant proportion of legal claims that are in fact resolvable by statutory provisions do not have any relevant 'corpus of legal cases' (particularly in jurisdictions of only modest size). It is not obvious that in such situations, ML extraction of the statutory provision that is most relevant to the problem would be better than a formal statute-based expert system that asked questions in order to identify the most relevant statutory provision, and then acts as a decision-support system to assist the human user to resolve issues of its interpretation. We argue therefore, that Ashley's paradigm is one of incomplete application, and that a statute-based decision support system may be a better paradigm for some problems and some use cases.

3.5 Aim to handle complexity

One main virtue of legal expert systems is the handling of complexity in a thorough way, to a level of comprehensiveness which is usually difficult to sustain by humans (except the most expert). This is rarely explicitly stated in articles, including ours, but it underlies the discussions of isomorphism[45], the need for declarative representations and similar matters. It is very easy to fail to take into account a definition of a term, or the implications of interconnection between sections, in complex legislation. One of the purposes of legal expert systems is to better ensure that users always take into account all relevant statutory provisions, and do not 'skip over' provisions which may appear to be irrelevant.

3.6 Users organisations should maintain their own knowledge-bases

The paradigm method of building expert systems assumed a division of tasks between a 'knowledge engineer' with experience in the methodology of building expert systems and the AI tools (particularly legal expert systems 'shells') used to build them, but with no necessary knowledge of a particular subject domain (eg law). However, competent knowledge engineers remained as scarce as unicorns (and probably as expensive). On this model, even though there may be many readily available 'domain experts' (lawyers expert in a specific field), if there are no 'knowledge engineers' to mine/ extract/ capture/ represent their knowledge in the formalism required by the shell, then no knowledge-base results. The resulting 'knowledge acquisition bottleneck' has always been the largest practical problem in the construction of legal expert/advisory systems. For these reasons, DataLex advocated from inception that 'the

[42] See Ashley, 2017, p. 13 and his summary of the goals of 'cognitive computing.'
[43] Ashley, 2017, p. 13.
[44] Ashley, 2017, p. 13.
[45] 'Isomorphism' in this context, refers to a one-to-one correspondence between a knowledge representation (eg the rules in a knowledge-base) and the legal source being represented (eg a the sections of a statute).

construction of legal expert systems is done directly by a lawyer without the use of a knowledge engineer intermediary.'[46]

If (as is likely for most applications for free legal advice providers) it therefore becomes desirable (or simply necessary) for lawyers to self-maintain the knowledge-bases for their own systems, this leads to considerable design constraints in both software and knowledge-bases, addressed in [3.7]-[3.14].

3.7 Use declarative knowledge representations where possible

Declarative knowledge representations state whatever legal knowledge is available about a field (eg the sections of a statute), but without any particular order of representation being required, or any order of processing specified. In contrast, procedural knowledge representations determine the order in which items in the knowledge representation will be processed. Knowledge representations can be a mixture of both approaches.

The use of declarative representations of legal content is usually called 'rule-based reasoning' or 'inferencing'. Typically, rules will be both backward and forward chaining by default. The goal of a consultation will be set by the user. The inferencing mechanism works by the execution of statements comprising a rule with the goal as its conclusion. Backward chaining is used to determine unknown facts. When a fact value is inferred or supplied by the user, evaluation of all rules which contain that fact in their conditions takes place by forward chaining. The process iterates until a value for the goal is determined.

Declarative representations are desirable where possible, because: isomorphism is easier to achieve; knowledge-base development is faster; knowledge representation is more transparent; and less maintenance of knowledge-bases is required. [47] Procedural representations are necessary for any type of document generation system, and for some problems requiring execution of steps in procedural answers.

Declarative knowledge representations of legislation appear to be assumed by Ashley, but with many reservations (which are clearly correct) about the limitations of the completeness of such representations because of such issues as the open texture of legislation.[48]

3.8 Isomorphic representations are desirable

'Isomorphism' is used in this context in the sense of 'creating a well defined correspondence between source documents and the representation of the information they contain [that is] used in the system' [49] – something close to a 'one-to-one correspondence'. Isomorphic representations are advantageous principally because 'a given source change can be related to a defined fragment of the knowledge base, and … this fragment can be removed from the knowledge base, altered and replaced with confidence that nothing else will be affected by the changes'.[50] Isomorphic representations of legislation are desirable (as far as possible) for creating, understanding and maintaining statute-based knowledge-bases, and for providing explanations of conclusions by reference to legislative provisions.[51]

---

[46] DataLex, 1987a, p. 6, 'Basic DataLex Design Criteria: 1 Direct construction by lawyers'
[47] DataLex, 1995: '4. The inference engine'
[48] Ashley, 2017, Ch. 2; He states (p. 8-10) that declarative representations are used by what he regards as the expert systems paradigm, and explains the roles of forward-chaining and backward-chaining rules.
[49] Bench-Capon and Coenen, 1992
[50] Bench-Capon and Coenen, 1992
[51] DataLex, 1993; 1995 – parts 5.5, 5.6; Bench-Capon, 1989, Karpf, 1989, Johnson and Mead, 1989, Bench-Capon and Forder, 1991, Bench-Capon and Coenen, 1991, Johnson and Mead, 1991, Bench-Capon and Coenen, 1992

Ashley agrees, arguing that 'the legal rule modeling language needs to support isomorphism',[52] noting particularly that this 'makes explanation more effective'.[53]

3.9 Quasi-natural-language knowledge-bases avoid repetitive coding

Quasi-natural-language knowledge representations express legislation (or other legal knowledge) in English-like sentences (or other human languages), with logical connectors (such as IF, THEN, ONLY-IF, NOT) also expressed in formalised natural language. They enable those writing knowledge-bases to avoid separate coding of questions and explanations, because questions, and all types of explanations, can then be generated automatically from the knowledge-base. This speeds both construction and maintenance. [54] Verbatim representation of predicates ('facts') within such sentences is also desirable.[55]

A representation in quasi-natural-language has these advantages for the maintenance of isomorphism: a structured English-like representation, with a limited range of logical operators, removes ambiguities better than natural English; 'textual baggage', detracting from isomorphism, is eliminated; transparency for validation is increased; transparency during use is increased; there is increased explanatory power; and there are benefits for application development, such as the relative ease with which domain experts can check rules written by others.[56]

This possibility seems to be overlooked by Ashley when he assumes that maintaining faithful isomorphic representations 'between statutory texts and implementing rules,' 'require maintaining multiple representations,' including so that 'decision aids such as textual excerpts from the statutory rules and links from commentary and cases can be linked into the program's logical explanation of a conclusion.'[57] With a quasi-natural language knowledge representation, coupled with automated hypertext links from its terms, there need only be one representation in order to achieve links to the statutory text, access to decision aids of different types, and generation of explanations of various types.

3.10 Propositional representation is enough for most tasks

The DataLex project developed a predicate calculus inferencing system (allowing multiple instances of variables) [58] which also included a quasi-natural language knowledge representation (Aide). It was a successful approach in that it increased the isomorphism of the knowledge representation, as well as being logically more powerful. However, it also increased the complexity of the dialogues with the user necessary to obtain results from the system, often in ways that could not be easily understood by the user. It was also not as easy for developers to understand what inferencing steps would be taken when an application ran, because (for example) a natural language representation of a section of an Act might be split by the parser of the knowledge-base into two or more rules (in Horn clause form). In this sense, the knowledge representation could be regarded as 'deceptively simple'.

In comparison, a YSH (or WYSH) knowledge-base was somewhat less isomorphic and less resembling natural (legislative) language, but had the advantage that the steps that would be involved in drawing inferences were more apparent to developers. Choice of tools usually involves trade-offs, and in this case our conclusion was that that most legal problems did not require multiple instantiations of variables, so a propositional representation like YSH/WYSH

---

[52] Ashley, 2017, p. 63.
[53] Ashley, 2017, p. 64.
[54] DataLex, 1995: '5.5.1. Isomorphism is facilitated by English-like representations'
[55] Poulin et al, 1993; DataLex, 1995: 5.5.1.
[56] See DataLex, 1995: '5.5. Isomorphism and the value of English-like representations' for support for each of these propositions.
[57] Ashley, 2017, p.64.
[58] DataLex, 2001

might be the better choice. Where multiple instantiations of variables were required, a tool (like Aide) which can handle predicate-calculus-like reasoning may be necessary.

Others have similarly concluded that for most legal expert system applications, forms of logic more complex than propositional logic will not be necessary.[59] Ashley points out the advantages of predicate logic, but his example of a Prolog representation of the British Nationality Act, while very informative, is also able to be represented in propositional logic, and it is not clear under what conditions he considers predicate logic essential.[60]

3.11 Semi-expert systems and users collaborate

The aim in building legal expert systems is not to build a 'robot lawyer', which simply extracts unproblematic facts from a user and then comes to a conclusion. Almost all systems require the user to provide some degree of interpretation of the questions asked, and the sources of law involved, requiring at least a minimal level of interpretative skills. The real model of a legal expert system[61] is therefore one of collaboration between a semi-expert computer system, and a semi-expert user, with control of a problem's resolution alternating between them.[62] The aim is to support decisions made by human users. The result is best described as a 'legal decision support system', rather than an 'expert system' or 'robot lawyer'.

This is the same fundamental approach as what Ashley describes as the 'second new paradigm for system development', 'cognitive computing': 'The operative unit of cognitive computing is neither the computer nor the human but rather the collaborating team of computer *and* human problem solver(s)' (emphasis in original).[63] Ashley cites 2013 research, but this aspect of the cognitive computing paradigm has deeper roots (including for DataLex since the 1980s). Rejection of what Leith describes as 'the robotisation of lawyers'[64] was from inception the basis of the DataLex approach.

However, depending on the context, users (the 'human problem solvers') may range from being trained lawyers who have some knowledge of the subject area of law involved in the system (but are not fully expert), to intelligent lay users or self-represented litigants. The design of the system, and the extent to which it relies on the user's expertise, will depend on which point on this spectrum the intended/likely users will be found. The implications of these models of 'AI and law' for the legal profession are still becoming apparent, and are often misunderstood.[65]

3.12 Inferencing is not enough for decision support

Access to legal sources and other forms of legal expertise is almost always necessary, except in the most trivial of legal expert systems, because interpretation issues cannot be eliminated from knowledge-bases. This means that inferencing systems cannot be 'closed': they must give users access to the legal sources on which interpretation is based. Because law is constantly changing (most notably, by the creation of new case law), if such access is to a limited set of resources ('closed' in another sense) it will be unsatisfactory. From a user

---

[59] For example, Bench-Capon, 1989.
[60] Ashley, 2017, pp. 47-49.
[61] DataLex, 1989c: 'Models of a legal expert system'
[62] DataLex, 1989c: 'The interactive user (interpretative?) model': 'The most useful general model by which we may conceptualise a legal expert system, seen from the perspective of the user of such a system, … seems to be that of an interaction between a semi-expert inferencing system and a semi-expert user/interpreter, with control over the course of the problem's solution alternating between the two parties to the interaction'.
[63] Ashley, 2017, pp. 12-14; see also p. 355: 'Cognitive computing has a similar goal, to achieve a kind of mass customization of legal advice, but shifts the focus to the human user by striving for an intelligent computer-human collaboration.'
[64] Leith, 2010, p. 5, and see footnote 18 where he mistakenly assumes that because we used the expression 'legal expert system' we are advocating what he is castigating.
[65] Greenleaf, 2017

perspective, inferencing systems must be as open as possible to all relevant legal resources, primary and secondary.

For developers, this has implications. There are four predominant models/paradigms of representing and manipulating legal information: text structures for search engines/ text retrieval; hypertexts; knowledge-bases for inferencing/advisory systems; and documents for generation. Each paradigm has well-recognised limitations and deficiencies when it is used on a 'stand-alone' basis for the computerisation of legal information.[66] Integration of inferencing, hypertext and text retrieval can overcome some of the limitations of each of these current paradigms of legal computerization, limitations which become very apparent when they are used on a 'stand-alone' basis.[67]

'In the cognitive computing paradigm', says Ashley, the embodied knowledge used by the system 'is embodied in the corpus of texts from which the program extracts candidate solutions or solution elements and ranks them in terms of their relevance to the problem'.[68]

3.13 Legal expertise can and should be captured by multiple means

Knowledge-bases are only one way of representing, storing and re-using legal expertise or knowledge, and not always the best way. Lawyers have long traditions of doing so through creating textbooks, law journal articles, citators, annotators, case reports and the like. All technologies for utilising legal information allow legal expertise to be stored and re-used, particularly hypertext (through links to content judged relevant) and text retrieval (through stored searches to find relevant content).[69] The 'relevant content' may be primary legal materials, but it may equally be commentary that is itself a means of capturing legal expertise.

3.14 Automate hypertext linking from knowledge-bases

The integration of hypertext and text retrieval with inferencing systems is facilitated by use of a (quasi) natural language knowledge representation [3.9]. Hypertext links and stored searches from dialogues and explanations – to definitions, cases, citators, commentary etc – can be both automated and hand-crafted, and increase the ability of users to both answer questions and to understand explanations.[70] These are means of combining different forms of captured expertise [3.11]. This can now be improved: [4.3].

3.15 Collaborative knowledge-bases can crowd-source development

Parts of knowledge-bases can be distributed across different websites, and invoked remotely as part of a consultation, using a closed wiki.[71] This is one means of distributing the maintenance and development costs of complex knowledge-bases, which may be necessary for some free access advisory systems. In recent decades, many models for such collaborative development of content, including open source software, open wikis and closed wikis: [4.4]. The development of legal expertise as a commons (a shared resource, usually free) both created and managed collaboratively, is a significant challenge not only in relation to knowledge-bases, but also in relation to more conventional ways of representing expertise such as online textbooks.[72]

---

[66] DataLex, 1992a, DataLex, 1995.
[67] DataLex, 1992a; DataLex, 1995: '1. 'Integrated' computerisation of law', and citations therein, particularly [Paquin et al 1991]
[68] Ashley, 2017, p.13.
[69] DataLex 1992a; DataLex, 1995: '1. 'Integrated' computerisation of law'; see also Greenleaf 2017, part IIB 'Representing Expertise', and Susskind and Susskind, 2015.
[70] DataLex, 1995 – part 5; DataLex, 1997 – parts 5, 8.
[71] DataLex, 1997 – part 7; DataLex, 2000.
[72] Greenleaf, 2017, parts IIB and IIC.

3.16 Summary: Conclusions relevant to free legal advice services

We can now attempt to summarise and apply what might be relevant from our conclusions arising from our experience with the DataLex project, to a free legal advice service. This is based on our assumptions of the likely limited financial and personnel resources of such a service [1.2], and that the 'AI and law' systems it could be expected to find useful are those that justify their answers at least in part in terms of the formal sources of law: [3.1]. These constraints will mean that only some types of 'AI and law' tools are suitable to their needs: [3.3].

First, looked at from the user perspective, which could be that of an employee of a free legal advice service, or perhaps one of its clients, what counts as a useful level of legal expertise is relative. A system may be valuable to a class of users even though it has a relatively low point at which it admits that a problem is beyond its expertise, and it may also serve as a method of triage [3.2]. In any event, it is not realistic to try to build legal expert systems that encapsulate all the knowledge necessary to answer user problems: [3.12]. The more realistic aim is to build decision support systems, in the use of which the program and the user in effect pool their knowledge/expertise to resolve a problem: [3.11]. Expertise can and should be represented and utilised by programs in many ways [3.13]. This means the knowledge-based system (the knowledge representation and the program) should not be 'closed': it must be integrated with text retrieval, hypertext and other tools which allow and assist the user to obtain access to whatever source materials are necessary to answer the parts of a problem dependent on the user's expertise: [3.15]. The result is an integrated decision-support system.

Second, looked at from the developer perspective, the key contextual factor is that user-organisations such as free legal advice services, will probably need to both develop and maintain their own knowledge-bases, as the only available domain experts [3.6]. The systems which non-technical legal domain experts are most likely to be able to develop and maintain are those which represent legal knowledge in a way which has a reasonably high level of isomorphism (one-to-one correspondence) with the legal sources on which it is based [3.8], where the representation is reasonably close to natural language [3.9], and where it is not necessary to prescribe the order(s) of the procedural steps necessary to reach a solution to a problem, but only to declare what legal knowledge is available, and leave it to the system to undertake the steps to apply that knowledge: [3.7].

Thirdly, correctly choosing the type of problem where 'AI and law' techniques are most likely to be appropriate is essential. Problem areas based on legislation, or procedural steps [3.4], and where there is complexity [3.5], will probably give the best results. Problems involving multiple instances of one factor increase logical difficulty: [3.10]. If it is organisationally possible to have multiple organisations collaborate to build and maintain a legal knowledge base, this may increase sustainability: [3.14].

Finally, it is worth noting that although Ashley, as a leading current proponent of the field of AI and law, might well regard our ambitions for system development as unnecessarily modest (or perhaps just the product of our constraining assumptions), there is little that is inconsistent between the 'legal decision support system' approach taken by the DataLex Project from the 1980s, and its underlying rationales, and the 'cognitive computing' paradigm advanced by Ashley thirty years later. Many aspects of the history of AI and law retain some relevance for the present and future, and ignoring them can waste time and resources.

4 Desirable improvements for sustainable legal advisory systems

Many new technical developments in recent years create a very different context for the creation of legal advisory systems than were present in the 'second wave' of AI and law (to 2001), in which the original DataLex Project was a participant. We outline here, at the conceptual level only, a number of desirable objectives that follow from our previous arguments, which were not generally achieved or feasible prior to 2001, but which are now more achievable. These include far better techniques for creating interfaces to inferencing systems and integrating them with other resources, and the availability of vast new bodies of primary and secondary legal materials, particularly for free access developments. There are also many new forms of AI techniques which can be used by the legal profession in various ways, but only some of these are relevant to the types of legal advisory systems we discuss (ie those that attempt to justify their advice or conclusions by reference to primary or secondary legal sources and reasoning based on them).

4.1 More supportive editing environments for knowledge-bases

The 'knowledge acquisition bottleneck' is probably still the main impediment to the practical implementation of legal advisory systems. This problem is exacerbated if it is necessary to have legal experts – who may have no computing skills beyond facility with a word processor – creating knowledge-bases directly, without any 'knowledge engineer' as an intermediary. We argue this is necessary for free legal advisory systems: [3.6]. Quasi-natural language knowledge representations reduce this problem, but more supporting development interfaces are needed to reduce it further.

Developments in the last decade or so have made it far easier to create a supportive editing environment to assist lawyers to do so. These include: visual programming, drag-and-drop programming and sophisticated integrated development environments. Developments in Natural Language Processing (NLP) and related AI algorithms also open up opportunities to explore other legal knowledge and semantic representation models. Further development of natural language and conversational interfaces[73] as well as more 'intelligent' interfaces may now be made possible by introducing a NLP layer on top of existing systems.[74]

4.2 Automated first-cut rule-bases for legislation

One of the most valuable enhancements to the development of legislation-based systems would be that when a lawyer specifies that a part of a piece of legislation is to be added to the knowledge-base, that part is then imported from a free access legislation source, and automatically converted, as far as is possible, into the syntax required by the knowledge representation. Such an 'automated first-cut rulebase' facility, would then be tested, and edited, by the author until it ran correctly. No matter how imperfect the conversion of the legislation, for a lawyer to be able to go immediately to some 'first draft' of part of a knowledge-base, with at least some of the rule syntax correct, would speed knowledge-base construction and reduce the knowledge acquisition bottleneck. This is an active area of research,[75] but Ashley does not report successful automated or semi-automated conversion of statutes into rulebases in his survey,[76] and assumes human encoding from inception.[77]

---

[73] For example, McTear, Callejas and Griol (2016) discuss the technologies behind conversational interfaces and how they can be applied in the context of smart devices

[74] Open source toolkits for the processing of natural language text such as OpenNLP <http://opennlp.apache.org/> make the development of such interfaces more feasible.

[75] For example Data61 'Regulation as a Platform' is a two year prototyping project that claims they are 'building an open platform based on a machine-readable version of current laws,' <http://www.data61.csiro.au/en/Our-Work/Future-Cities/Optimising-service-delivery/RaaP>. They have made rule-sets available <https://2017.hackerspace.govhack.org/dataset/regulation-platform>. The project aims to produce 'first cut' rulebases which must be edited by experts before use: see T. McGregor and N. Patel 'A digital prototype to help reduce the complexity

4.3 Automated linking to more sophisticated legal sources

As we have argued [3.12], inferencing systems are of limited value in isolation: they must be integrated with open-ended sources of legal information, usually via hypertext or text retrieval, so as to enable users of such systems to make legal interpretative decisions when it is necessary for them to do so because the system cannot. The availability of such links from many different aspects of legal advisory systems (dialogues, reports etc) is essential from the user perspective. The automation of such linking is essential from the developer perspective.

Compared with twenty years ago, there are now vastly more legal sources available for free access (essential for free legal advisory services), and they are of much greater sophistication than provision only of legislation and judicial decisions (citators,[78] commentary[79] etc). It should therefore now be possible to create more sophisticated and valuable linking from inferencing systems to sources enabling interpretation, and also for these sources to have far more sophisticated links between these sources than was possible previously. For example, an inferencing system may create a link to the most significant case (at the time the knowledge-base was written) on an issue that the user must decide. However, if the case linked to also provides access to a citator linking to all subsequent cases that have considered it, then the system is to some extent 'self updating'. The DataLex software, integrated with AustLII and the LawCite citator, now achieves this. Ashley considers that the US Ravel system 'introduces legal researchers to the virtues of citation networks on a large scale', and 'will be important for cognitive computing'.[80]

4.4 Better tools for collaborative development of knowledge-bases

Since the 1990s, there have been very great advances in both the tools for, and widespread experience in, collaborative development of software (free and open source software – FOSS) and in collaborative development of open content (for example, Wikipedia). These developments have also resulted in shared resources (public domain or commons). These developments are part of the collaborative embodiment of expertise.[81] In law, there have been many failed attempts to develop sharing of expertise in text,[82] but there are some notable successes including CanLII Connect (case annotation) and Wex (legal dictionary) from the LII (Cornell). These two rare successes originate from legal information institutes.

Lawyers providing different free access advisory services have collaborated to produce substantial textbooks of expert commentary, using the AustLII Communities platform of a closed wiki.[83] Examples include an online collaboration of seventy contributors from legal aid and other services in Australia's Northern Territory to produce its comprehensive 'Law Handbook'.[84] Multiple collaborators have also produced the Australian Capital Territory Environmental Law Handbook.[85] It is therefore feasible to consider that different legal services could also collaborate to produce shared advisory tools, if provided with a

---

of regulation', *ABCB Connect*, undated < http://www.abcb.gov.au/Connect/Articles/2017/05/15/A-digital-prototype-to-help-reduce-the-complexity-of-regulation>.

[76] Ashley, 2017, Ch 2 'Modeling Statutory Reasoning'.

[77] Ashley, 2017: see for example, pp. 60-62 on development of business process compliance rules; see also p 355, referring to Waterman's 'manual analysis of legal texts'.

[78] AustLII's LawCite citator <http://www.austlii.edu.au/lawcite/> indexes over 5.5 M cases from across the common law world and back to the thirteenth century, plus journal articles and treaties; see Mowbray, Chung and Greenleaf, 2016.

[79] For example, O'Neill and Peisah *Capacity and the Law* (2nd Ed, 2017, AustLII Communities) <http://austlii.community/wiki/Books/CapacityAndTheLaw/>

[80] Ashley, 2017, pp. 353-54.

[81] Greenleaf, 2017

[82] R Ambrogi 'The failure of crowdsouring in law (So far, at least)', *Law Sites blog*, 10 August 2015. Notable failures include JurisPedia.

[83] AustLII Communities < http://austlii.community/foswiki/Main/WebHome>.

[84] The *Northern Territory Law Handbook* < http://austlii.community/foswiki/NTLawHbk/NTLawHandbook>..

[85] <http://austlii.community/foswiki/ACTEnvLawHbk/ACTEnvLawHandbook>

collaborative platform such as AustLII Communities or CanLII Connects. As yet, collaborative legal knowledge-base development has not been applied to a significant practical extent[86] although this has been advocated.[87]

4.5 Platforms to assist development and maintenance of free advisory systems

Free legal advice services will usually only have limited access to computing expertise, and therefore limited expertise to host and maintain the operation of online legal inferencing/advisory systems: [1.2]. There is a potential role for a provider of a shared computing facility, like a legal information institute, to host and maintain both the development tools ('shells') and applications developed using them (whether individually or collaboratively). Development of the knowledge-bases would still be left to the lawyers of the legal advisory services, and updating (the other sense of 'maintenance') must remain the responsibility of such domain experts.

4.6 Data-driven and data-oriented feedback tools

Access to greater legal content and improvements in data analytics allow for the development of more contextualised feedback during consultation. For example, during a consultation involving a specific legislative section, related case-law and other relevant resources can be presented. This will result in a better overall user experience and will help in obtaining better input in relation to open-textured questions. Enhancements in web technologies may lead to more responsive and interactive experience for the user of AI systems. For example, websockets and other new web technologies enable the creation of web applications that make use of real-time, persistent, and bi-directional messaging,[88] so as to improve user experience by automatically making relevant context information available during expert system consultations. Many of these tools are available in open source versions, and are therefore available to free advisory services, provided the costs of expertise in deploying them can be found.

4.7 Renewed attention to transparency and ethical operation of legal analytics

The increasing use of 'big data' analytics, including in systems developed for the legal profession and for government administration means that there must be renewed attention to ensuring that the algorithms employed in AI applications to law are understandable to both those who are using the systems, and those people whose lives may be adversely affected by them. In relation to legal advisory systems, there is already a substantial body of research on the explanation functions, transparency and isomorphism of such systems, and research on such functions needs to continue and be improved. The differences between such systems and those AI applications to law which are based on correlations, or other reasoning with weak relationships to formal legal sources [3.3], needs to be stressed, and the systems critiqued. Free legal advisory systems (or legal information institutes) are less likely to have a strong proprietary interest in the knowledge-bases they employ than are commercial law firms, and are therefore more likely to be willing to open them to scrutiny.

5 Conclusion: Potential for free access/free advice collaborations

The argument presented in this paper is that research and development in the previous waves of interest in 'AI and law', particularly in the area of legal expert systems, suggest many conclusions and techniques which are still relevant and valuable. This is particularly so if the

[86] DataLex, 2000 demonstrated how this could be achieved, in an academic setting, and is believed to be the first such application.
[87] Susskind and Susskind, 2015; Greenleaf, 2017
[88] Pimentel and Nickerson (2012), for example, discuss the performance of websockets and other polling techniques for real-time display of content.

focus is on how providers of free legal advice, with necessarily limited means, can make use of AI developments.

We have referred throughout this paper to the specific examples of AustLII, and the DataLex project. Free legal advice services (by whatever names) in other countries are likely to be considering how they can best use 'AI and law' technologies to improve their services to clients. Many countries have legal information institutes (LIIs)[89] and other types of providers of free access legal information. The conclusions drawn from experience in this paper do not have any particular geographical location, and we hope they will be of value to others considering similar issues.

While legal expert system technologies can be readily applied to any area of law with some superficial success, they are not a panacea, and in many areas of law may add little to what available human expertise, textbooks and checklists can provide. The factor that makes the difference is to find those relatively rare applications of the technology that really do justify the costs involved in developing them, where the result is a legal advisory system that changes people's lives for the better. This will usually require domain experts – such as free legal advice centres – to identify the application areas where they would like to be involved in development. The key challenge is for LIIs and free legal advice service to find achievable applications that meet the needs of clients and add value to decision-making that cannot otherwise be achieved at equivalent cost or quality. The key measure of success of such systems must be what the users of these applications have to say about their value (and the objective factor of their extent of use, not what their developers say, measures which Leith claims were routinely ignored in the previous wave of 'AI and law'.[90] If in the past, 'very few programs ever were tested in any meaningful way or even made available for scientific assessment by others,'[91] this also points toward the value of either free access or open content programs and applications.

[89] See Free Access to Law Movement (FALM) <http://www.falm.info/> and its membership.
[90] Leith, 2010, pp. 6-7.
[91] Leith, 2010. p. 7.

6 References

In this reference list, all publications which were part of the DataLex Project 1985-2001 are grouped as authored by 'DataLex' for convenience, and the full author details then given. A history and bibliography of the DataLex Project 1985-2001, is on SSRN at <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3095897>, and a version with functioning links is at <http://www2.austlii.edu.au/~graham/expert_systems.html>.

(Ashley, 2017) Ashley, K *Artificial Intelligence and Legal Analytics* (Cambridge, 2017)

(Capper and Susskind, 1988) Capper, P and Susskind, R *Latent Damage Law: The Expert System* (Butterworths, 1988)

(Bench-Capon, 1989) Bench-Capon, T 'Deep models, normative reasoning, and legal expert systems' (1989) *Proc. ICIAL*, 37.

(Bench-Capon et al, 2012) Bench-Capon, T et al 'A history of AI and Law in 50 papers: 25 years of the international conference on AI and Law' (2012) 20(3) *Artificial Intelligence and Law*, 215–319.

(DataLex, 2001) Allen, R, Chung, P, Mowbray, A, and Greenleaf, G (2001) 'AustLII's Aide – Natural Language Legislative Rulebases' at *Proc. International Conference on Artificial Intelligence & Law* (ICIAL) 2001 (SSRN); also 'AustLII's Aide – Implementing Rulebase Systems' [2001] CompLRes 32, as presented at 5th Law via Internet Conference, 2001. This paper concerns the use of predicate logic in the Aide software.

(DataLex, 2000) Allen, R, Greenleaf, G, Austin, D, Chung, P, and Mowbray, A, (2000), 'With a Wysh and a Prayer: An Experiment in Cooperative Development of Legal Knowledge Bases', *The Journal of Information, Law and Technology* (JILT), vol 2 (SSRN); shorter version also published in (1999) *Proceedings of the 7th International Conference on Artificial Intelligence and Law* (7 ICAIL), Oslo, Norway (ACM Press).

(DataLex, 1997b) Greenleaf, G, Mowbray A, King, G, Cant, S and Chung, P, (1997) 'More than Wyshful Thinking: AustLII's Legal Inferencing Via the World Wide Web' Proceedings of the *6th International Conference on Artificial Intelligence and Law*, Melbourne, ACM Press, New York (SSRN). This is the main paper setting out how the DataLex Project was migrated to the web (and the AustLII context) and the advantages this provided.

(DataLex, 1997a) Mowbray, A, Greenleaf, G, King, G and Cant, S, *Wysh Developer's Manual* [1997] COL 1, Australasian Legal Information Institute, Sydney (SSRN).

(DataLex, 1997) Greenleaf, G, Mowbray, A and King, G 'New Legal Services via the Web - AustLII's Research on Legal Inferencing' (Chapter 5 of 'New Directions in Law Via the Internet: The AustLII Papers', Journal of Information, Law and Technology (JILT), No. 2, 1997). This was the first paper published about WYSH. (SSRN).

(DataLex, 1995a) Mowbray, A 'Info-One alternative system demonstration' (consultancy project) – AIRS and HYPE application.

(DataLex, 1995) Greenleaf, G, Mowbray, A and van Dijk, P (1995) 'Representing and Using Legal Knowledge in Integrated Decision Support Systems: Datalex Workstations' *Artificial Intelligence and Law*, vol 3, nos 1-2, pp 97-124 (SSRN). This is the major paper summing up the DataLex Project work to 1995 (ie pre-web and pre-AustLII).

(DataLex, 1992a) Greenleaf, G, Mowbray, A and Tyree A, 'The Datalex Legal Workstation – Integrating Tools for Lawyers', (1992) 3(2) *Journal of Law and Information Science* 219 -240 (SSRN) (also in *Proc. Third Int. Conf. on Artificial Intelligence and Law*, ACM Press, 1991).

(DataLex, 1992) Greenleaf, G, 'The Privacy Workstation' (1992) 6 *International Yearbook of Law Computers and Technology*, pp177 – 196 (SSRN).

(DataLex, 1991a) Mowbray, A 'DataLex Workstation system' – software (Unix, DOS/Windows, Macintosh) (integration of YSH, AIRS and HYPE into one package).

(DataLex, 1991) Mowbray, A 'YSH (y-shell) inferencing system' – software (Unix).

(DataLex, 1989d) Tyree, A *Expert Systems in Law* (Prentice Hall (Aust), 1989).

(DataLex, 1989c) Greenleaf, Graham, 'Legal Expert Systems - Robot Lawyers? (An Introduction to Knowledge-Based Applications to Law)' *Proc. Australian Legal Convention*, Darling Harbour, Sydney, August 1989 (SSRN).

(DataLex, 1989b) Mowbray, A and Greenleaf, G 'LAWS OF AUSTRALIA Defamation demonstration' – HYPE and AIRS application (Unix).

(DataLex, 1989a) Mowbray, A 'HYPE hypertext system' – software (Unix).

(DataLex, 1988c) Mowbray, A, Greenleaf, G, Tyree, A & Lewis, D (1988), 'Teaching Lawyers Online Retrieval: The AIRS Training System.', pp. 96-102.

(DataLex, 1988b) Greenleaf, G, Mowbray, A and Lewis, D *Australasian Legal Information Retrieval Handbook*, Butterworths, 1988

(DataLex, 1988a) Tyree, A, Greenleaf, G and Mowbray, A (1988) 'Legal Reasoning: the Problem of Precedent' in J Gero & R Stanton (Eds) *Artificial Intelligence Developments & Applications*, North-Holland, Amsterdam, 1988, pgs 231–247 (not on SSRN).

(DataLex, 1988) Greenleaf, G, Tyree, A and Mowbray, A (1988) 'Communications Aspects of Legal Expert Systems — Incorporating Them in Shells for Lawyers' *Proceedings of the National Conference on Law, Computers & Artificial Intelligence* University of Exeter, United Kingdom, November 1988, 10pgs (SSRN).

(DataLex, 1987c) Greenleaf, G and Griffith, P 'COPYRITA – copyright law inferencing system' (LES application).

(DataLex, 1987b) Greenleaf, G, Mowbray, A and Tyree, A (1987) 'Legal Expert Systems: Words, Words Words...?' *Yearbook of Law Computers & Technology* Vol.3, 1987, 119–136, Butterworths (UK) (SSRN). This is a much-expanded version of the Boston ICAIL paper, including details of the COPYRITA, INTEST and FINDER applications, and the LES, AIRS and PANNDA software.

(DataLex, 1987a) Greenleaf, G, Mowbray, A and Tyree, A (1987) 'Expert Systems in Law: The DataLex Project' (*Proc. First Int. Conf. on Artificial Intelligence and Law*, ACM Press, 1987) (SSRN) The first paper about the DataLex Project, and the first paper presented (with a demo of LES) at the first ICAIL Conference, Boston, 1987.

(DataLex, 1985c) Tyree, A 'FINDER – finding cases inferencing system' – FINDER application.

(DataLex, 1985b) Mowbray, A 'INTEST – Intestacy legal inferencing system' – LES application.

(DataLex, 1985a) Tyree, A 'PANNDA (precedent analysis by nearest-neighbour discriminant analysis) inferencing system' – software.

(Feigenbaum and McCorduck, 1983) Feigenbaum, E A and McCorduck, P *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World* (Addison Wesley Publishing, 1983).

(Greenleaf, 2017) Greenleaf, G 'Review Essay: Technology and the professions: Utopian and dystopian futures' (2017) 40(1) *UNSW Law Journal* 302-321.

(Johnson and Mead, 1989) Johnson, P and Mead, D 'Legislative Expert systems' (unpublished), SoftLaw, Canberra.

(Johnson and Mead, 1989b) Johnson, P and Mead, D 'Natural language - An appropriate knowledge representation scheme for legislative expert systems' (unpublished), SoftLaw, Canberra.

(Johnson and Mead, 1991) Johnson, P and Mead, D 1991 'Legislative knowledge base systems for public administration' *Proc. 3rd ICAIL* ACM Press.

(Karpf, 1989) Karpf, J 'Quality Assurance of Legal Expert Systems', (1989) *Jurimatics* No 8, Copenhagen Business School.

(Leith, 2010) Leith, P 'The rise and fall of the legal expert system' (2010) 1(1) *European Journal of Law and Technology.*

(McTear, Callejas and Griol, 2016) McTear, M, Callejas, Z and Griol, D *The Conversational Interface: Talking to smart devices* (Springer) 2016.

(Michie and Johnson, 1984) Michie, D and Johnson, R *The Creative Computer* (Pelican, UK) 1984.

(Mowbray, Chung and Greenleaf, 2016) Mowbray, A, Chung, P and Greenleaf, G 'A Free Access, Automated Law Citator with International Scope: The LawCite Project' *European Journal of Law and Technology* (EJLT) (2016) Vol 7 No 3.

(Paliwala, 2016) Paliwala, A (2016) Rediscovering artificial intelligence and law: an inadequate jurisprudence?, *International Review of Law, Computers & Technology*, 30:3, 107-114

(Pimentel and Nickerson, 2012) Pimentel, V and Nickerson B G 'Communicating and Displaying Real-Time Data with WebSocket' *IEEE Internet Computing*, 16 (July/August 2012), 45-53

(Pollack, 1992) Pollack, A '"Fifth Generation" Became Japan's Lost Generation', The New York Times (online), 5 June 1992.

(Sergot and Kowalski, 1986) Sergot, M J, Sadri, F, Kowalski, R A, Kriwaczek, F, Hammond, P, and Cory, H T (1986) 'The British Nationality Act as a logic program'. *Commun. ACM* 29, 5 (May 1986), 370-386.

(Susskind, 1986) Susskind, R 'Expert systems in law: a jurisprudential approach to artificial intelligence and legal reasoning' (1986) 49 *Modern Law Review* 168.

(Susskind and Susskind, 2015) Susskind, R and Susskind, D *The Future of the Professions: How Technology Will Transform the Work of Human Experts* (OUP, UK) 2015

(Waterman et al, 1986) Waterman, D A et al 'Expert systems for legal decision making' *Proc. Second Aust. conf. on Applications of Expert Systems*, Institute of Technology, Sydney, 1986.