Legal knowledge based systems
JURIX '95
Telecommunication and AI & Law

The Foundation for Legal Knowledge Systems
Editors:
J.C. Hage
T.J.M. Bench-Capon
M.J. Cohen
H.J. van den Herik

Carolus Grütters, A simulation model for the Dutch asylum procedure: a different approach of law and AI, in: J.C. Hage, T.J.M. Bench-Capon, M.J. Cohen, H.J. van den Herik (eds.), *Legal knowledge based systems JURIX '95: Telecommunication and AI & Law*, Lelystad: Koninklijke Vermande, 1995, 53-64, ISBN 90 5458 252 9.

More information about the JURIX foundation and its activities can be obtained by contacting the JURIX secretariat:

Mr. C.N.J. de Vey Mestdagh
University of Groningen, Faculty of Law
Oude Kijk in 't Jatstraat 26
P.O. Box 716
9700 AS Groningen
Tel: +31 50 3635790/5433
Fax: +31 50 3635603
Email: sesam@rechten.rug.nl

 http://jurix.bsk.utwente.nl/

A SIMULATION MODEL FOR THE DUTCH ASYLUM PROCEDURE: A DIFFERENT APPROACH OF LAW AND AI

Carolus Grütters
Law and Information Technology
University of Nijmegen, P.O. Box 9049, 6500 KK Nijmegen, The Netherlands
email: carolus@jur.kun.nl

Abstract

Research in the field of Law and Artificial Intelligence (AI) is aimed at the development of legal knowledge-based systems (LKBSs) that solve legal problems. The direction of this research has been greatly influenced by two presuppositions. The first is that the legal problems to be solved relate to *individual* cases. The second one is the focus on *output.* However, if we start from a different point of view, changing the interpretation of legal problems, and focus on legal systems instead of individual cases and throughput instead of output, AI-and-Law research reveals new promising applications. This paper presents such a different view and describes its opportunities illustrated by a prototype of a simulation model of the Dutch asylum procedure.

1 Introduction

AI and Law research is aimed at the development of LKBSs that are capable of solving legal problems. In trying to do so, major obstacles were identified and the aims of research shifted from complete systems towards more modest contributions on different parts of the original goal (de Wildt *et al.*, 1990, p. 114; Grütters, 1995a), particularly with regard to legal reasoning. In that perspective Law and AI has been defined as: *the design and analysis of computational methods for performing or supporting legal reasoning tasks.*[1] Although most of the research in the area of AI and Law falls within the scope of this definition, there is an interesting counterpart of research which has no explicit relation to legal reasoning as such.

The first and main reason why legal reasoning is thought to be the major part of the definition follows from the presupposition that LKBSs should solve (complex) legal problems, i.e., reason with specific aspects of either individual or classes of *cases.* Systems dealing with *individual* cases simulate the activities of a judge (van den Herik, 1991) by performing legal reasoning tasks whereas those dealing with *classes* of cases, support the activities of legal drafting. Although the research on knowledge-based support for legal drafting is quite different (Bench-Capon, 1992), it is also based on the presupposition that (questions about) *cases* should be solved. Questions regarding the consistency and completeness of regulations or sets of legal rules are then in fact still considering cases, though at a higher level of abstraction than the individual case.

A second reason for the emphasis on legal reasoning in the definition of Law and AI has to do with the focus on the *output* of the problem-solving processes. In other words, LKBSs generate judgements on individual cases or on classes of cases, i.e., legislation. This classical approach of Law and AI might be described with the metaphor of the physician who has to diagnose and decide on the therapy.

Leaving this commitment to cases and output aside, implies a new interpretation of the legal problem to be solved by LKBSs. That interpretation does not concern solving cases but gathering *insight* into complex legal structures. In doing so, systems can be built that simulate legal procedures, reveal information about the behaviour of legal systems, particularly with respect to the effects of changes within a procedure, giving less

[1] See the introduction of the editors in Prakken *et al.* (1994, p. 3).

attention to the output and more to the throughput. This other approach of Law and AI might be described with the metaphor of a supervising unit which monitors and analyses all activities.

2 The problem of change of legislation

Law, or a legal system, is never constant and is continuously changing, either directly by changes of the rules, or indirectly by change of their interpretation. In other words, law is not static but dynamic. For that very reason, various authors have raised the serious problem of maintenance of LKBSs (Weusten, 1989; Oskamp, 1990; Bratly *et al*., 1991; Bench-Capon and Coenen, 1991).

Legal systems, considered as sets of rules, also tend to grow. Authorities trying to solve social conflicts by (changing) legislation are more likely to increase than to decrease the number of rules. Any increase of number, however, results in an increase of complexity. The main problem of this complexity is a lack of overview, and an insufficient insight into both internal and external effects. From that perspective it is understandable that authorities desire more information about the impact of their measures, before these measures are put into practice.

2.1 Monitoring the external effects of change

EXPERTISZE (Svensson, 1993) and TRACS (den Haan, 1993) are relevant applications concerning the simulation of effects of (changes in) legislation. However, these effects are only external effects: information about the change in (types of) individual cases and their output resulting from certain changes in legislation. The same goes for MISISZ (Lathouwer, 1993) which concerns a comparative analysis tool for social-security systems, monitoring external effects.

In spite of the importance of this information on external effects, there is no information about the internal effects: how will certain changes affect the procedure itself, e.g., the practice or the workload of the administration dealing with the cases. In addition to the importance of knowing what effects will be caused by certain measures, it is vital to know if and how these measures affect the procedure itself.

2.2 Monitoring the internal effects of change

Outside the legal domain many simulation tools are available which can monitor the actual throughput in a system and simulate its behaviour. Within the legal domain, however, it has not been an interesting item to know, for example, whether a change in a procedure influences the workload of the administration or the total handling or cycle time. This is probably due to the fact that, for years, the legal system was there, waiting for a call for justice. In due time all calls were dealt with. At the moment, however, the existence of backlogs, the limited capacity, and the need for priorities in legal procedures, underline the importance of the question whether the structure of certain, if not all, legal domains is adequate to deal with all the calls. To answer that question we need information about the behaviour of a legal system, i.e., insight into its structure.

3 A system-dynamic approach

The attempt to obtain insight into a system entails dealing with its dynamics. One way to approach a domain in such a way has been explored by Forrester (1961, 1968) and has been developed further under the name of System Dynamics, closely related to Operational Research and Systems Thinking (Richardson, 1994; Wolstenholme, 1990).

Using a system-dynamic approach of a certain process implies the construction of a dynamic model which simulates the behaviour of that process in time. Such a model is built in order to obtain some insight into complex structures and to observe the behaviour of a system, or parts of it, i.e., to see whether a simulation run of such a model shows tendencies of growth, stabilisation or decline. These are precisely the items policy makers are interested in, when considering a change of policy or legislation.

This approach is well-known in other domains than law, such as the handling of goods, business procedures, industry and environmental studies on food supply, or population growth. A system-dynamic approach applied to a legal domain is rather new. However, the first results show an interesting field of research, and practical applications closely linked to other efforts in the area of AI and Law. By its structure there is, or seems to be, a resemblance between legal procedures and business or administrative processes. The point is that until recently (Grütters, 1995b), there has not been any investigation to examine whether law is an interesting domain for this type of approach.

Whatever the reasons may be, the main characteristics of the legal domain play an important role here, namely the separation of powers and the existence of checks and balances. The consequence of this is that cooperation between judiciary, legislator and executive is often confused with dependence, and is seen as incompatible with this separation. This results in a situation in which relevant information is not, or only partially, available to those who might need that information. One might say that this is a consequence of preventing a Big Brother effect, or just the result of bureaucracy. The fact remains, however, that insight into the relations and interdependencies of different actions in legal procedures is at least problematic. Following the observation that legal procedures resemble administrative processes, it is worthwhile to investigate what differences exist and whether legal constraints create obstacles.

In my opinion there are two main differences. The first difference is that legal procedures as such are far more important than the result. Justice is characterized by its procedure, by its way of dealing with people, not just by the outcome of a procedure. In other words, a just procedure is the result of justice, whereas administrative or business procedures are only necessary obstacles to achieve a certain result. Otherwise, justice could be done more efficiently by organizing a lottery or by throwing dices. Therefore, legal constraints are part of the system, and in that perspective the outcome of a just and legal procedure is irrelevant in terms of justice. Thus, legal constraints are, in a way, comparable to security or quality checks in a production line, and may be regarded as vital parts, which should not be avoided but handled with care. Nevertheless, this does not mean that legal procedures cannot be improved both in terms of efficiency and of justice.

The second difference is the possibility of following different routes generating similar outcomes, or following the same route with different results, which complicates the attempt to obtain an overview of the whole procedure. Lawyers are not interested in this type of roadmap overview. They are trained to analyse a specific given route until a certain point, and will advise on how to proceed from the viewpoint of their client's interest. The process of analysing all possible routes from start to finish is therefore, in general, unknown to lawyers. By contrast, this type of information is very relevant to policy makers and those involved in drafting legislation. However, these officials are almost always lawyers, especially within the Ministry of Justice.

3.1 Asylum law: a legal domain

The Dutch asylum procedure has been selected as a test domain for a system-dynamic approach. This choice is based on a number of factors, such as the presence of expertise both on refugee law and System Dynamics within the University of Nijmegen. In addition, the Dutch asylum procedure has been altered quite drastically in 1993, and again in 1994, aiming at the simplification of the procedure and reduction of the workload, i.e., elimination of backlogs. The main motives behind the changes in the legislation were:

- an increasing number of requests,
- the existence of backlogs and long cycle times throughout the asylum procedure, and
- the desire to reduce these substantially.

Despite the changes in legislation, the Dutch asylum procedure still is complex and requires a large number of officials, varying from officials from several offices within the Ministry of Justice to interpreters, lawyers, counsellors, judges and their respective administrative staffs.

Since no tools were available to simulate the impact of the proposed changes, the main part of the public debate about these changes concentrated on their political feasibility; not on their effectiveness or efficiency. The only facts available consisted of data concerning the number of requests, i.e., the input, and the number of decisions, i.e., the output. What happened in between was largely unknown, and any substantial insight into, or overview, of the throughput in the whole procedure was lacking. In other words, the actual debate was based on opinions how the procedure should be, instead of how it would be.

3.2 The asylum procedure

In short, the asylum procedure starts with a request from an asylum seeker and ends whenever a final decision has been made by the administration, implying that the request has been granted or denied. At first sight, this seems a simple process from start to finish. However, the asylum procedure involves two different but parallel procedures. The first one deals with the actual request for admittance to the country and recognition as a refugee under the Dutch Aliens Act and the Treaty of Geneva on refugees of 1951. The second one concerns the lodging of the asylum seeker during the time he is awaiting any decision on his initial request. Both procedures are closely interconnected. For example, a certain decision on an initial request might imply that another type of lodging is required and that the asylum seeker has to move, either out of the country or towards another housing facility.

After a negative decision in first instance, the administrative part of the procedure provides a review, possibly followed by an appeal. Moreover, when the administration - at the Ministry of Justice - decides to deny the request and wants to expel the asylum seeker, it is possible that the asylum seeker starts a separate legal procedure aiming at the prevention of his expulsion based on the assumption that the outcome of a yet undecided review or appeal of his request may be to his advantage. To complete the procedure's complexity, it has to be stated that there are at least two different ways to start this procedure and that there are at least four different ways of granting the request.

It will not take a great deal of imagination to see that there is a large variety of different routes between start and finish of the asylum procedure. Combined with the fact that all parts of the procedure have their own demands, in terms of time and capacity, we are faced with a complex system. It is difficult to say, then, what impact certain changes will have on parts of the procedure. In order to decrease this complexity and to create a certain amount of overview and insight, a system-dynamic approach has been examined.

4 SMAP: a Simulation Model of the Asylum Procedure

4.1 Investigation

The Simulation Model of the Asylum Procedure (SMAP) is composed of several consecutive stages. First, we listed all parties involved and their separate responsibilities for the decision moments in the procedure. Based on that list, we depicted all possible - and not just likely - routes between start and finish along these decision moments.

Secondly, we made an inventory of the problems by consulting experts at the Ministry of Justice, the judiciary, lawyers, officers working in the centers where the asylum seekers are temporarily lodged, and experts on asylum law. The main points on that list concerned the queues at several points - or decision moments - in the procedure, capacity problems and the idea that, whatever measure was taken, the workload on the administration remained at least the same. In other words, there are bottlenecks at different permanent or temporary places, and it is not clear what, if any, effect a certain measure will have and when. This also means that it is not clear what measure has to be taken in order to achieve a certain proposed effect.

Thirdly, we tried to gather data on the asylum procedure. Although that seemed to be the easiest part, it became the most tedious one. The idea of providing researchers with production and capacity figures, data on handling time and the like, appeared to be very frightening, notably to the Ministry of Justice. First we had to convince the officials that the results might be of interest to them, and consequently that the data we wanted were mainly meant to verify our prototype. Apart from the fact that the Ministry of Justice has easy access to the data, which could also have been gathered at various other places, though with more effort, it then appeared that the Ministry of Justice had only recently started to collect these detailed data. Finally, we mention that the Ministry's anxiety was merely based on a lack of information, which was exactly one of the reasons why we had started the research. Gradually, however, its reserved attitude changed into cooperation.

4.2 Mapping

We arranged all the available data on a Power Macintosh using, in first instance, the software package Stella II® and, later, Ithink®. In order to achieve and to keep as much overview as possible, we defined several layers. The top layer (Figure 1) entails a conceptual or mental

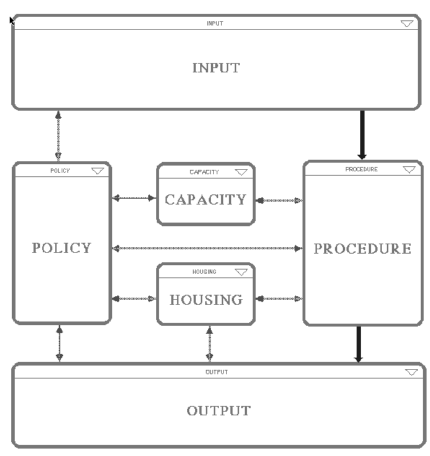

Figure 1: The conceptual model (top level)

model of the system and consists of several frames representing separate entities within the procedure: policy, capacity, housing, input, legal procedure and output. These frames are connected in two ways: first, by thick arrows representing actual flows (of files) and second, by thin arrows representing information flows. Figure 1 shows the flow of files from the input frame into and through the procedure frame and then into the output frame. These actual flows are influenced by information flows of the other frames. For instance, a significant increase of asylum seekers may lead to a certain change in policy consisting of changes of capacity, housing or procedure. These changes then influence the actual flow of files through the procedure and into the output frame.

Below this top level, a second level was defined which shows, in more detail, the various sub-frames corresponding with the separate phases of the procedure in each of the previous frames. Figure 2 shows a part of this level, focusing on the file flow in the administrative part of the procedure, leaving out the information flows. Within each of these sub-frames the actual procedural part is elaborated, using the various building blocks available: pipelines, stocks, converters and connectors (see Figures 3 and 4).

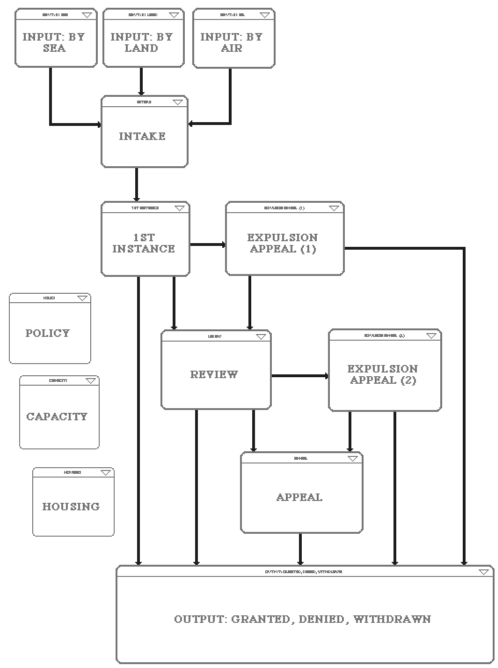

Figure 2: The flow of files in the administrative part

Since any sub-frame is connected to other sub-frames, there are several in- and outflows, shown as clouds. Within the sub-frame the actual flows (of files) are depicted using pipelines with taps. These taps indicate the possibility of regulating the flow through the pipes, while the flows go in and out certain stocks.

The open box (in Figure 3) represents a reservoir. The striped box represents a conveyor in which it takes some time to flow through. This very important feature enables the introduction of (variable) delays and transit times. Since we are interested in the behaviour of the system in time, it is vital to model these delays and other related time functions. Another feature of this conveyor is the possibility to define limits on the actual input or throughput.

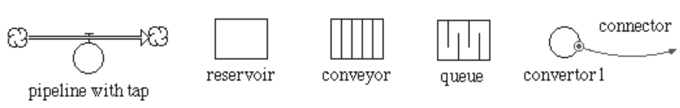

Figure 3: Building blocks of the actual procedural part

The half striped box represents a queue, the status of which depends on the situation beyond the queue. If, for instance, a conveyor would be defined after a queue with a certain limited input per time unit, the queue would be filled at the moment that the input of the queue exceeds its output, i.e., the input limit of the conveyor. In the same way, the queue will empty whenever the input of the queue is below its possible output, i.e., the input limit of the conveyor.

Converters, depicted as small circles, represent mathematical operations which may act on taps or other converters. And finally, connectors, depicted as bent arrows, represent the connections between the building blocks.

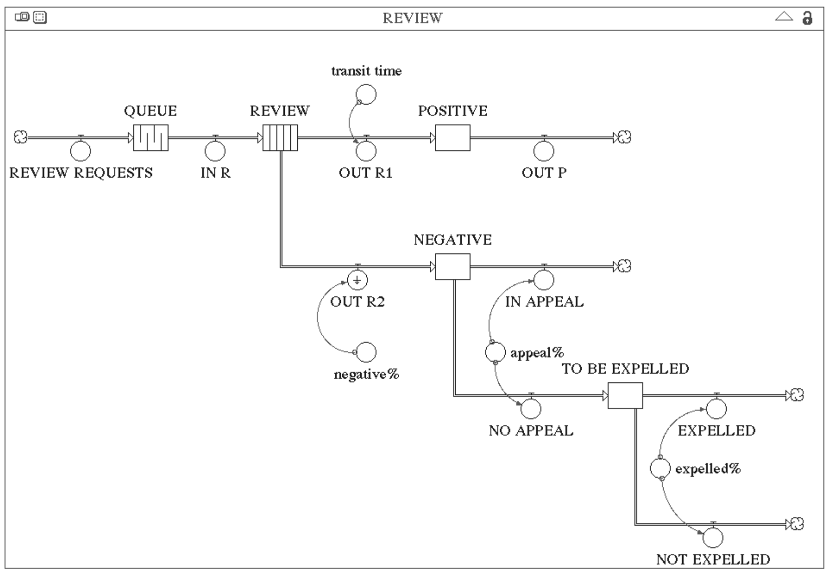

Figure 4: The connection between building blocks

4.3 Tuning

After modeling the procedure and the initial filling of all elements with data, the next phase of tuning could be started. At this stage the research aims at verifying the structure of the model and at tuning the values and formulas used in different building blocks. Since there is a known input (asylum requests) and a known output (granted requests, denied requests and withdrawals) per time unit, it is possible to examine whether running the model generates correct or acceptable figures for all sectors somewhere in the procedure that match the data from other sources.

The main part of the official data consists of month-data. We therefore entered all time related data as month-data and selected a total run-time of ten years, since the average cycle time for one file was somewhere between one and two years, and a limited number had cycle times of four or five years. In order to prevent the model handling the month-data as one quantity, we selected a time unit of months and $\Delta t = 0.25$, meaning that the month-data would be divided into four separate quantities and that the whole model would be computed four times per time unit.

Since the first runs of the model - which only had an input, a throughput and an output frame - generated data which were close (i.e., about 90%) to the data of registered practice, we continued expanding and tuning the model. At the same time it has to be emphasized that, although the model generates exact numbers during the simulation in all building blocks animated or graphically displayed in graphs or tables, the value of the model does not lie in the prediction range of exact numbers. Nevertheless, the model's strength lies in the possibility to detect tendencies as decline, growth or stabilization. Precisely this type of information is of great value to policy makers: they are likely to be more interested in knowing that a certain situation of continuous growth will occur after a certain period of time than that the exact number of N will be present at stage S at time T. The same goes, of course, for any other tendency.

Another important aspect of this tuning concerns feed-back loops and the selection of delays. Feed-back loops are mainly defined within the frames that generate information flows, such as capacity figures, which influence the actual production within parts of the procedure. Delays are particularly used wherever, for example, the moving of files from one point to another takes more time than just one time unit, as the map suggests.

5 Analysis of simulation runs

Although the tuning of the model is not finished yet, our present research focuses on the analysis of simulation runs. This analysis serves different purposes. In addition to verifying the model, it enables the gathering of insight into the behaviour of the system as a whole as it reflects the present situation in the asylum procedure in general terms.

At this point a critical remark has to be made. Although one would expect policy makers to be interested in general information such as tendencies and overviews, research in the related area of micro-simulation has shown *that policy makers are often particularly interested in those "details", rather than in global distributive changes* (Lathouwer, 1993, p. 117). I would like to add that this preferred interest for details might be the result of a mixture of a lack of overview, an inability to look beyond present political balances, and a practical quantity of self-interest. Besides, most changes of policy are about details and from a political point of view this is a way of surviving. Bearing in mind this pessimistic caution, it may also be the case that the tools available for policy makers are just too detailed.

5.1 Scenarios

Due to a number of factors, such as capacity shortage, simulation runs show bottlenecks and patterns in the procedure of, for instance, continuous growth. Most of the time this is undesired. There are mainly two ways to change this, or to detect a set of measures which will lead to a different, more acceptable pattern.

The first method is based on trial and error, meaning that different scenarios are formulated which contain various types of changes in the model. These scenarios are more or less based on intuition and experience with the simulation model. Although such a method may come up with interesting data based on a comparison of different scenarios and their results, it still remains incomplete. Provided that the procedure is in fact complex and not surveyable without the help of tools - as is the case in the Asylum procedure - it will remain uncertain whether there is another scenario that will approximate the intended behaviour of the system in a more desirable way.

The second method involves the use of AI techniques aiming at an analysis of all possible routes, combined with a complete overview of possible alterations within the system that could lead to a certain predefined behaviour of the system. This type of optimization research goes further than a regular sensitivity analysis. It entails the search for different optimal structures of a procedure that uses intelligent tools, bearing in mind that the user must be able to define the limits of certain parameters. This latter feature should prevent a combinatorial explosion and at the same time force the user to formulate the limits of both the procedure and the practical possibilities.

Valente (1995, pp. 133-135) discusses models of legal tasks within the context of a legislation-oriented typology of legal problems based on a suite of problem types described by Breuker (1994). In addition, Valente elaborates two models for assessment and planning, i.e., normative assessment and agent legal planning. His planning (or reconstruction) model seems to be useful within our context of the use of system dynamics, although the internal effects on the administration are not mentioned. However, the models that might be even more interesting, are the models of prediction (postdiction) and monitoring. Whenever such an approach would be possible, and the internal effects would be taken into account, a legal procedure might be evaluated in a more exhaustive way.

5.2 Insight

In addition to the use of scenarios, task models and the possible search for other structures and optimal behaviour, it remains that a system-dynamic approach is in the very first place meant to learn how a certain system works and how it behaves in time. Strictly theoretically, a system-dynamic approach is already successful whenever the construction of a model leads to insight into previously unknown relations and dependencies.

However, insight alone is not enough in practice. Knowing how a traffic jam occurs and in what way certain measures interact in decreasing or increasing queues, one also has to choose between all the options available aiming at a “better” handling of the traffic, whatever meaning “better” is given in this context.

6 Conclusions

The use of a system-dynamic approach to a legal domain is quite different from any other legal AI research. According to my analysis this is due to the fact that most research has concentrated on static systems dealing with cases which produce custom-made solutions. In terms of metaphors this view is based on the concept that the activities of a legal problem solver resemble the work of a traffic policeman, determining the best route for each car that comes along. The system-dynamic approach is, however, based on the metaphor of a traffic map designer or planner, who considers in what way a given amount of traffic can be regulated within a certain period of time.

Whether one of these two different approaches should be preferred is not an interesting question: it depends. The point is that this other approach is also possible and may result in more knowledge about how law works, and particularly how legal procedures can be made transparent, and justice improved.

The main differences between these two types of approaches concern the focusing on internal relations within a system and taking into account the factor time. In addition, a system-dynamic approach does not imply costly and time-consuming gathering of data on individual cases.

The modelling of legal procedures by use of a system-dynamic approach, possibly extended with the use of certain AI techniques, creates the possibility of monitoring the effects of change, particularly internal effects, expressed in terms of behaviour of the system. This type of ex-ante evaluation by means of simulating legal procedures might support policy makers and those involved in drafting legislation, without taking the risk of experimenting in real practice at the expense of vulnerable people, such as asylum seekers.

Software References

The model described in this paper is developed using Stella II® and Ithink®, both registered trademarks of High Performance Systems Inc. Hanover, NH, USA.

References

Bench-Capon, T.J.M., and Coenen, F.P. (1990). Practical application of knowledge based systems to law: the crucial role of maintenance. *Legal knowledge based systems: Aims for research and development*, pp. 5-17. Proceedings of the JURIX ’90 conference, Koninklijke Vermande, Lelystad.

Bench-Capon, T.J.M. (1992). Support for Policy Makers: Prospects for Knowledge Based Systems. *Legal knowledge based systems: Information Technology and Law*, pp. 42-49. Proceedings of the JURIX ’92 conference, Koninklijke Vermande, Lelystad.

Bratley, P., Frémont, J., Mackaay, E., and Poulin, D. (1991). Coping with change. *Proceedings of the Third International Conference on Artificial Intelligence and Law*, pp. 69-76. ACM, New York.

Breuker, J. (1994). A suite of problem types. *CommonKads Library for Expertise Modelling* (eds. J. Breuker and W. van de Velde). IOS Press, Amsterdam.

Forrester, J.W. (1961). *Industrial Dynamics*. MIT Press, Cambridge MA.

Forrester, J.W. (1968). *Principles of Systems*. MIT Press, Cambridge MA.

Grütters, C.A.F.M. (1995a). Visions on IT and Law. *Law, Computers and Artificial Intelligence*, Vol. 4, No.1, pp. 113-126.

Grütters, C.A.F.M. (1995b). The simulation of legal procedures. *Proceedings of the 1995 European Simulation Multiconference*, 433-438. SCS, Belgium.

Haan, N. den (1993). Towards support tools for drafting legislation. *Legal knowledge based systems: Intelligent tools for drafting legislation, computer-supported comparison of law*, pp. 23-30. Proceedings of the JURIX ’93 conference. Koninklijke Vermande, Lelystad.

Herik, H.J. van den (1991). *Kunnen computers rechtspreken?* Inaugural speech (in Dutch), Gouda Quint, Arnhem.

Lathouwer, L. De (1993). Adequacy of Unemployment Schemes: an Application of Microsimulation in Comparative Analysis. *Legal knowledge based systems: Intelligent tools for drafting legislation, computer-supported comparison of law*, pp. 115-124. Proceedings of the JURIX ’93 conference, Koninklijke Vermande, Lelystad.

Oskamp, A. (1990). *Het ontwikkelen van juridische expertsystemen (een theoretische beschouwing)*. Doctoral thesis (in Dutch). Kluwer, Deventer.

Prakken, H., Muntjewerff, A.J., Soeteman, A., and Winkels, R.G.F. (eds.) (1994). *Legal knowledge based systems: The Relation with Legal Theory*. Proceedings of the JURIX ’94 conference, Koninklijke Vermande, Lelystad.

Richardson, G.P. (1994). Systems Thinkers, Systems Thinking. *System Dynamics Review*, Vol. 10, No. 2-3.

Svensson, J.S. (1993). *Kennisgebaseerde microsimulatie: een nieuwe methode voor het bepalen van sociaal-economische gevolgen van wet- en regelgeving in de sociale zekerheid*. Doctoral thesis (in Dutch). Universiteit Twente, Enschede.

Valente, A. de Souza Mello (1995). *Legal Knowledge Engineering; A modelling approach*. Doctoral thesis. IOS Press, Amsterdam.

Weusten, M.C.M. (1989). Maintenance of knowledge in Advisory Systems on Legal Questions. *Preproceedings of the third international conference on “Logica, Informatica, Diritto* (ed. A.A. Martino), pp. 853-862. CNR, Florence.

Wildt, J. de, Schmidt, A., Quast, J., Herik, J. van den, Kralingen, R. van, and Visser, P. (1990). The evolution of research aims. *Legal knowledge based systems: Aims for research and development*, pp. 106-117. Proceedings of the JURIX ’90 conference, Koninklijke Vermande, Lelystad.

Wolstenholme, E.F. (1990). *System Enquiry: A System Dynamics Approach*. Wiley, Chichester.