



7th BILETA Conference Information Technology and Legal Education: Towards 2000

9th & 10th April 1992

**Sponsored by
Context Ltd; Linklaters & Paines; Needham & James**

A Prototype Expert System For Making Bail Recommendations

Patricia Hassett

Syracuse University College of Law

Syracuse, New York, USA

Keywords: because statements - card sorts - common sense rules - decision model - decision trees - discretion - expert Systems - expert system shell - field test - heuristic rules - interface - interlocking domains - matrices - Outcome monitoring - prototype - 'what if' capacity - XiPlus

Abstract: The state has multiple interests in the quality of bail decisions. Previous papers (Hassett, 1991, 1990) explored the issue of whether computer technology, specifically expert systems technology, could contribute to improving the quality of pretrial release / detention decisions. Subsequently a prototype expert system which makes bail recommendations was developed using the XiPlus expert system shell. This paper documents the process by which the prototype was developed, identifies the capacities and limitations of the prototype, and suggests how the bail expert system project might be carried forward.

1. Introduction

The Bail Act attempts to balance the conflicting interests. However, a combination of circumstances (including the difficulties inherent in assessing risks such as flight and re-offending) complicates the application of the Act and makes the quality of release/detention decisions hard to monitor and to control. As a result, the Bail Act may not be achieving an appropriate balance.

The expert system is a type of computer software which attempts to mimic human thinking. This project attempted to harness the expert system technology to the bail decision process to see if the quality of the bail decisions could be improved. Because computer systems are usually time-consuming and expensive to develop, the computability of the problem and the feasibility of the project are usually first tested by trying to build a prototype (or sample) of the full scale system.

2. The Process Used to Develop the Prototype

2.1 The Partnership:

The prototype was a collaborative project. I brought legal expertise and a limited technological

literacy. Nigel Payne, from Andersen Consulting, brought technological expertise and a limited legal literacy. We worked together very satisfactorily. The project took about four months with the two of us working more or less full time.

2.2 The Selection of an Expert System Shell:

I entered the project with no familiarity with expert system shells. Discussions with expert system aficionados had produced a variety of recommendations, including Crystal and XiPlus. Andersen Consulting, which had a copy of XiPlus and some familiarity with it, recommended that we consider using it. Nigel Payne explored its facilities and found that it had a number of features that would help us achieve our goals.

2.3 The Goals of the System:

The idea of an expert system for bail decisions was prompted by a variety of concerns about bail decision procedures and outcomes.

One concern is that the bail decision maker has just moments to make a complicated decision; the pressures in that situation work against the goal of a well-reasoned decision. So, one goal is to provide an intelligent assistant which will enable bail decision makers to make rational decisions quickly and under pressure.

A second concern is that similar cases can result in quite different decisions. By providing a common set of consistently applied rules, the expert system can promote the goal of having like cases treated alike.

A third concern is that decisions are not explained, either orally or in writing. The absence of explanations creates problems on several fronts. Without an explanation, even a proper decision may seem irrational to the accused (and to the accused's lawyer, family and friends). Also, without an explanation, the accused may not know what additional facts would be useful in getting a more favourable bail decision.

Explanations of judicial decisions are useful for other reasons. They provide a basis for appellate review of the decisions in individual cases and a basis for an evolving understanding of the bail decision area in groups of cases.

2.4 The Prototype Issue:

The implication of a prototype is that its successful completion provides the basis for believing that the projected larger system is achievable. However, the portion of the problem selected for prototyping should be representative of the larger problem. Therefore, the selection of the issue to be addressed in the prototype is very important.

The bail decision involves three basic questions: Will the accused reappear as needed during the criminal proceedings? Will the accused re-offend pending the disposition of the criminal proceedings? Will the accused interfere with the administration of justice (by, say, intimidating or bribing witnesses)? We elected to prototype the question of failure to appear.

In order to further reduce the scope of the prototype, we decided to limit the prototype to one class of offenses: summary offenses punishable by imprisonment up to six months. Summary offenses are relatively minor offenses which are triable only in magistrates' court and without a jury. We excluded summary offenses punishable by a fine only on the ground that a person not subject to a custodial penalty would normally not be detained pending disposition of the case.

2.5 The Decision Rules: Formulation Methods:

2.5.1 The Expert: The Credibility Factor. The normal procedure for an expert system which attempts to replicate human thinking is to identify an expert and to have the expert explain the decision process. The use of the acknowledged expert is very desirable because it contributes to the user perception that the solutions proffered by the system are credible.

This procedure could not be followed in the bail domain for at least two reasons. First, bail decisions are not monitored or evaluated. As a result, the bail experts, if any, are not identified as such. Second, bail decision makers are not told what process to use to make the assessment. As a result, the assessment is not made by reference to a common process.

In the absence of an acknowledged expert and an accepted uniform process for making bail decisions, we elected to use an alternate route to credibility. We would construct the system with common sense rules. I would select the rules, not by reference to the views of a single expert, but by reference to a common sense belief (derived both from the bail literature and from experience) in their relevance. The credibility of the system would flow not from the user's acceptance of the expertise of the rule generator, but from the user's acceptance of the legitimacy of the rules themselves.

2.5.2 Legislative and Judicial Rules: Library Research. The user is able to make an assessment of the rules' legitimacy because the bail domain differs from other legal expert system domains in at least two important ways.

First, the bail domain rules do not involve sophisticated legal analysis. The rules are incorporated in a reasonably clear statute which says an accused is to be released pending the disposition of charges unless there are substantial grounds to believe he will fail to appear. The difficulties inherent in this formulation relate to factual assessments rather than to legal analysis: what factual circumstances should count as substantial grounds to believe?

The factual analysis called for by the statute has not been complicated by a body of judicial opinions which have encrusted the statutory rule with judicial modifications, exceptions, or other intricacies. As a result, the intervention of an acknowledged expert is not necessary to ensure that the system represents an authoritative way through the legislative and judicial rules.

Second, some legal expert system domains interlock in significant ways with other domains. When an interlock occurs, the system must either incorporate rules from the related domain, or tell the user that the system is unable to solve the problem. The interlocking means that the user may need help from an expert to identify the interlocking issues. The bail domain, however, is quite independent of other domains. As a result, special expertise is not necessary.

2.5.3 Common Sense Rules: Interviews and Observation. The first step in generating the rules was to explain the bail system to the computer expert. The explanation took the form of recorded interviews in which the Bail Act factors (details of the offence and of the accused) were related to the risk of non-appearance. So, for example, the risk of non-appearance is higher if the likely penalty upon conviction is custodial than if the likely penalty is non-custodial. Or the risk of non-appearance is lower if the accused has strong and well-established family and community ties than if the accused has no ties.

Repeatedly during the interviews, reference was made to the unavailability of, or the unreliability of, information relevant to the bail decision, and to the speed and chaos characteristic of the decision process. The computer expert grew increasingly uncertain that the interviews were generating reliable common sense rules. So the interviews were augmented by court visits at which actual cases were observed and subsequently discussed with court personal.

During the initial interviews, each factor was related directly to the risk of non-appearance. After all of the relevant factors were identified and related to the risk of non-appearance, the relationships among the factors needed to be addressed.

2.5.4 Relationships Among Facts, and Among Rules: Because Statements, Card Sorts, Matrices and Decision Trees. The identification of the relationships between factors and the risk of non-appearance was only a first step. In a case where all of the factors suggested that the risk of non-appearance would be low, the accused's release without conditions would be the easy conclusion. In the more usual case, however, some factors would suggest a low risk of non-appearance and others would suggest a high risk of non-appearance. As in real life, the system needed a way to reconcile the inconsistent suggestions and to reach a solution. We used four practical tools to achieve the reconciliation: because statements, the card sort, decision trees, and matrices.

Because Statements. After the relationships between factors and risk of non-appearance were identified, the nature of the relationship was described in a 'because statement'. For example, the fact that the accused has regular employment is relevant (according to the Bail Act and to experience) to the likelihood that the accused will fail to appear. The nature of the relationship is that the risk of non-appearance is lower if the accused is regularly employed than if the accused is not regularly employed.

The 'because statement' explains why the relationship is thought to exist. So, in the case of regular employment, the 'because statement' might say: regular employment reduces the risk of non-appearance because the accused will be unwilling to give up the benefits of employment (such as income to assist in a defense, potential character witnesses, a personal circumstance likely to be taken into account upon conviction on the issue of a custodial versus a non-custodial penalty).

The 'because statements' were helpful in showing that the risk of non-appearance involved two subsidiary questions: what is the risk of an intentional non-appearance (as a result of flight) and what is the risk of an inadvertent non-appearance (as a result, say, of a misunderstanding as to the re-appearance date). The risk of intentional non-appearance had a relationship to some facts because those facts showed whether the accused had something to lose by flight; the risk of inadvertent non-appearance had a relationship to some facts because those facts showed whether regular communication with the accused could be maintained. Some facts had a relationship to both intentional and inadvertent risk of non-appearance; regular employment, for example, gives an accused something to lose by flight and provides a means of communication with an accused.

Card Sorts. We used the card sort to identify relationships among factors. For example, at the initial interview, the existence of regular employment and the existence of permanent accommodation were each said to decrease the risk of non-appearance. The initial interviews did not address the case where an accused had no regular employment but did have a permanent address.

For the card sort, each factor, its impact on the risk of non-appearance, and its because statement were entered on an index card. The index cards were then sorted into two groups: one contained factors relevant to intentional non-appearance and the other contained factors relevant to unintentional non-appearance. For example, employment (as indicated above) was thought to reduce the risk of intentional non-appearance because the accused would be reluctant to lose the benefits of employment. Accommodation was thought to reduce the risk of unintentional non-appearance because the accused could be contacted by mail or visit, and so be less likely to miss an appearance because of a communication difficulty.

Where the factor was relevant to both intentional and inadvertent failure to appear, duplicate cards were prepared. Strong family ties, for example, could make an accused unlikely to flee, and could also provide a means for communicating with the accused.

Within each group, a variety of sortings were tried. Gradually, the factors were sorted into primary, secondary, and even tertiary factors. For example, both the likelihood of a custodial sentence and a significant prior record started out on an apparently equal footing as factors likely to increase the risk of non-appearance. The card sort resulted in the likelihood of a custodial sentence being a secondary factor, relevant in deciding whether the accused presented a significant risk of intentional non-appearance. The significant prior record became a tertiary factor, relevant to the determination of whether the accused was likely to receive a custodial sentence.

Decision Tree. Once the relationships among primary, secondary and tertiary factors were identified, a decision tree could be constructed. The decision tree displayed the relationships among all of the factors.

So, for example, the risk of non-appearance was determined by reference both to the risk of intentional non-appearance and to the risk of inadvertent non-appearance. The risk of intentional non-appearance in turn was determined by reference to factors such as the likelihood of a custodial sentence which in turn was determined by reference to factors such as the presence of a significant prior record. Where a given factor can have more than one value, the decision tree shows the impact of each value on any other factors which depend upon the given factor.

Matrices. Where one factor was dependant upon two or more other factors, matrices were used to integrate the impact of the independent factors on the dependent factor.

For example, as mentioned above, the ultimate issue was whether the accused presented a substantial risk of non-appearance. That issue was dependant upon the risk of intentional non-appearance (which could have one of three values: high, medium, or low) and upon the risk of unintentional non-appearance (which could also have one of three values: high, medium, or low). The matrix for integrating the possible values of the two factors is shown below:

3. Then The Risk of Non-Appearance Is

	High	High	High	High
If the				
Intentional	Medium	High	Medium	Medium
1 Failure to				
Appear	Low	High	Medium	Low
Value is				

High Medium Low

2. And If the Inadvertent Failure to Appear Value is

This matrix assigns the value 'high' to the risk of non-appearance if either independent value is 'high'. It assigns the value 'low' to that risk if both independent value are 'low'. It assigns a value of 'medium' to the risk of non-appearance if both independent values are 'medium' or if one is 'medium' and the other is 'low'.

The values assigned via this matrix to the risk of non-appearance, and via other matrices to other factors, are estimated values rather than scientifically determined values. Experience (in the form of an analysis of the levels of non-appearance found in cases where the expert system has been used) may suggest that the assigned values need to be adjusted upwards or downwards to achieve the lowest level of detentions or of conditioned releases possible, while also achieving an acceptably low level of non-appearances.

2.6 Facts for the Decision Rules:

The variety found in the human condition suggests that an enormous number of facts could be relevant to the issue of non-appearance. If the expert system attempted to anticipate and make

provision for all potentially relevant facts, it would probably not succeed in including all such facts, and it would require the users to collect and enter a paralysing quantity of facts. The system would quickly be seen as defective in the former case and as oppressively burdensome in the latter.

As a result, we decided not to have the system analyze every possible factor. Rather, the system focuses upon the factors identified as relevant by the Bail Act, and likely to be relevant in every case. Such factors include the details of the offense (its seriousness, and the likely sentence) and the details of the accused (prior record, employment status, accommodation arrangements, and the like).

In many cases, the statutory factors will be the only factors and the expert system will, as a result, assist in the evaluation of all factors in the case. In other cases, the existence of unique but relevant factors will mean that the expert system can offer assistance only with respect to the statutory factors; in those cases, the bail decision maker can start with the expert system evaluation of the statutory factors but will then have to consider whether the additional factors increase, decrease, or leave unchanged, the risk assessment produced by the system.

2.7 Relationship between Recommendations and Decisions:

The prototype makes a recommendation, not a decision. The distinction is important because it underscores the fact that the system's product is intended to begin rather than to end the evaluation of an accused's pretrial liberty. The system stops short of a decision for several reasons.

The most important reason has to do with the nature of the bail decision: it will have a profound impact on the life of each accused and so should not be made mechanically. The system can only be as good as the facts entered and the rules used to process the facts. In each case, the facts, the analysis, and the recommendation need to be reviewed to ensure that they are appropriate under the circumstances. Where the recommendation is not appropriate, the decision maker has the opportunity, indeed the duty, to depart from the recommendation in making the decision.

An important part of the review involves hearing the arguments on behalf of the accused. Even if the facts are accurate, the accused may shed light on the appropriateness of the recommendation. Given the traditional importance accorded the opportunity to be heard, a system which failed to provide that opportunity would have difficulty in gaining support from any of the constituent groups in the criminal justice system.

Another reason, for treating the system's product as a recommendation rather than a decision, relates to the comprehensiveness of the system. As indicated earlier, a system which took account of all possibly relevant facts would require impossibly burdensome fact gathering and data entry. So the rationale of the system is to gather and process data on factors likely to be relevant in all cases. This data produces the recommendation. At the court hearing, the bail decision maker can take any other facts into account during the evaluation stage between the recommendation and the decision.

As a result, by stopping short of making a decision, the system is able both to preserve the element of judicial discretion that contributes to the fairness of the justice system, and to respond to the wide range of unique facts that can be expected.

3. The Capacities and Limitations of the Prototype

3.1 Recommendation Explanations:

An important feature of the system is its capacity to provide an explanation for each recommendation. The explanation consists of two parts: a list of the facts taken into account and a

decision tree that shows both the rules which have been triggered by the facts and the way that the rules work together to produce the recommendation.

The fact portion of the explanation is a recapitulation of all of the facts taken into account in reaching the recommendation. In any given case, some of the facts will involve data obtained from the accused or from the prosecution. Whether the accused has a fixed address, or a prior record, are facts which may be obtained in this way.

In most cases, however, information about one or more important facts may not be available. When actual data is not available, the system uses default data. So, when the facts about an accused's accommodation or prior record are not available, the system defaults will apply: for accommodation, the default is 'no accommodation' and for prior record, the default is 'no prior record'. Some of the default values assigned are favourable to an accused ('no prior record') and some are not ('no accommodation').

The premise in assigning default values was that the accused was in the best position to supply favourable life style and social information; so, information not forthcoming would be presumed to be unfavourable in terms of the bail decision. The state, on the other, would be in the best position to supply information about the crime and the accused's prior record; so, the absence of information in these areas would be presumed to be favourable to the accused in terms of the bail decision.

The recapitulation of both actual and default facts enables the fact gatherer (and the data enterer, if different) to make sure that the correct facts have been entered. The recapitulation also enables the accused and the state to be satisfied that the facts are accurate and to provide any additional facts that maybe relevant. If a review of the recapitulation reveals either errors in the actual facts or reliance upon default facts when actual facts are available, the system permits the user to enter the corrections and then reevaluates the facts to produce a new recommendation.

The recapitulation may also show that one or more important facts are simply not taken into account by the system. For example, if one of two otherwise similar accuseds were undergoing a course of medical care such as regular dialysis, the accused needing the medical treatment would probably present a lower risk of flight than the other. The prototype, however, does not take medical care per se into account and so would not distinguish between the two. This fact would be one that the decision maker would need to take into account, along with the recommendation, in making the bail decision. At present, the explanation does not help the decision maker to evaluate additional facts (by, say, suggesting the kinds of facts that would warrant departing from the recommendation), but that is a feature which could be added at a later stage.

The explanation also lists every rule that was relevant to reaching the recommendation. This feature is intended to ensure that the accused is informed about how the decision is reached. This explanation is important in the general sense that explanations are useful to dispel misunderstanding and the appearance of arbitrariness. The explanation is also important in the particular sense that it enables an accused to see what information or kinds of information might be useful to support a request for a particular bail decision.

The explanation also provides a basis for judicial review of the bail decision. At the moment, explanations are normally not given because the press of business usually forces the decision maker to proceed directly to the next case. As a result, an appellate court has no satisfactory basis in most cases for deciding whether or not the bail decision maker selected the correct, or a correct, decision. The system would cure this problem by issuing the recommendation and the explanations simultaneously and instantaneously. The explanation would thus be available for appeal purposes. Of course, if the rules used in the explanation are rational and used in appropriate ways to reach recommendations, the bail recommendation will pass scrutiny.

A more difficult problem, which the system does not yet address, is how to handle supplementary facts which seem to warrant departure from the recommendation. Here, the press of business may prevent the decision maker from explaining fully and clearly why the departure from the recommendation is necessary. One possibility is to provide an appropriately designed space on the recommendation/explanation report for the decision maker to enter the supplementary facts, to characterize them as increasing or decreasing the risk involved, and to indicate what decision has been made. The explanation report would then provide a complete report of the decision and would be readily available for use on appeal.

An additional advantage of the explanation¹ particularly where it contains explanations of departures from the recommendations, relates to the development of the rules for bail decision. Because so few bail decisions are explained, no collective wisdom on bail issues has emerged. The explanations of the departures would enable the system's rules to be improved and refined over time, particularly if (as is suggested below) the decisions were regularly monitored and evaluated against the occurrences of non-appearance or re-offending.

3.2 'What if' capacity:

The XiPlus expert system shell has a "what if" capacity the user (or an assistant) to enter new data quickly and at the last minute, and to generate a new recommendation instantly.

This feature is important in the bail context because the bail decision process is frequently complicated by the late arrival of new or corrected information. The 'what if' capacity means that the new or corrected information can be taken into account without the necessity for delaying the proceedings. Perhaps equally important is the fact that the new or corrected information is neither over-valued nor under-valued during the stress and confusion of this last minute reprocessing of the available data.

The 'what if' capacity can be useful in another way. Sometimes the bail decision process is complicated by uncertainties about the reliability of some of the information. The 'what if' capacity means that the decision maker can process the data both under the worst case scenario and under the best case scenario. A comparison of the two recommendations and their explanations will quickly show whether the uncertain information is important to the decision. If the recommendation is shown to be totally uninfluenced by the uncertain information, the decision maker can proceed without delay and without committing resources to the resolution of the uncertainty. A decision to proceed might also follow from the finding that the impact of the uncertain information is slight. If, on the other hand, the uncertainty is shown to be very important to the decision, the delay and resource commitment involved in resolving the uncertainty will be shown to be justified.

3.3 User Interface:

The user interface is one of the less satisfactory features of the prototype at its present stage of development. Nigel Payne was not able to spend as much time as he would have liked on the interface. Even so, his efforts have produced a much more user friendly interface than is provided in basic XiPlus. He created some screens which permit the required data to be entered without too much difficulty.

One of the pluses of the interface is that it edits out questions that are not necessary to the recommendation. One of the minuses of the interface is that it does not permit the user to back up to change the replies given to earlier questions. New or altered information can only be entered by starting all over again, or by using the 'what if' capacity after the initial recommendation has been generated.

3.4 Outcome Monitoring:

One important goal of the larger bail expert system project is to monitor the success of the system. This means ascertaining what percentage of the releases (whether conditional or unconditional) were inappropriate. The difficulty is that the failure of a release does not mean that the decision to release was wrong: a correct conclusion that the risk is low does not become an incorrect conclusion just because the risk materializes in the form of a non-appearance or a re-offending.

The prototype does not address this goal. If the system is developed beyond the prototype stage, some form of monitoring will be necessary in order to adjust the rules to improve the quality of the recommendations. One of the major challenges ahead is designing an appropriate process for the needed monitoring.

3.5 Judicial Acceptance:

An important concern about the prototype is that it has not yet been tested with bail decision makers. Judges and magistrates have traditionally been concerned with preserving judicial independence and with ensuring an individualized justice. These concerns could cause bail decision makers to be wary of any mechanistic approach to the exercise of their duties.

At the same time, judges and magistrates are concerned with the inappropriate disparities in treatment accorded to substantially similar cases. They are very aware that independent decisions are not judicial, and individualized decisions are not justice, if the decisions reached seem arbitrary and capricious.

As a result, some balance must be reached between independence and individualization, on the one hand, and consistent treatment, on the other. This system offers the possibility of an appropriate balance by creating a two step process: first, common rules are used to process common facts and to produce consistent recommendations; next, the decision maker is given the opportunity to adopt, modify or reject the recommendation according to the unique facts in the case.

Hopefully, this approach to balance will recommend itself to the decision makers. Of course, without judicial acceptance of the system and its features, the project can not go forward.

4. The Way Forward

The next steps for the project are clear, but not necessarily easy.

The prototype should be tested in the field. The testing would have at least two goals: first, to see whether the system makes recommendations that are seen as credible to the participants in the criminal justice system; and, second, to see whether the expert system can find acceptance as an intelligent assistant.

If the response to the field testing is encouraging, the prototype could be expanded to include an assessment of the risk of re-offending. The system's ability to deal successfully with this issue would be particularly timely given the publicity now accorded to 'bail bandits'.

Another round of field testing, this time of the expanded prototype, could include a monitoring of the relationship between the system's recommendations and the incidence of flight and re-offending. Presumably, the system would have a future only if it were to achieve flight and re-offending levels lower than (or at least equal to) the levels which occur with decisions made without the expert system.

Finally, a process for routinely re-evaluating and adjusting the rules in the expert system, in the light of experience, would have to be devised. Ideally, this process could involve at least some of the bail decision makers who use the system. Such involvement would permit the system to be enhanced by

their insights and would permit them to be confident that the system will enhance rather than supplant their judicial function.

References

Hassett, P. "Problems in Selecting Effective Computer Technology for Use at the Bail Stage of the Criminal Justice System". Proceeding of the 2nd National Conference on Law, Computers and Artificial Intelligence(1990).

Hassett, P. "An Expert System for Improving the Pretrial Release/Detention Decision". Pre-Proceedings of the 6th Annual BILETA Conference (1991).

Author

Patricia Hassett

Professor of Law

Syracuse University College of Law Syracuse, New York

USA

and The Lord Chancellor's Advisory Committee on Legal Education and Conduct

London