# Development of a KBS for Personal Financial Planning Guided by Pragmatic KADS

Susanne Dirks
IBM Ireland
Dublin

John K.C. Kingston[1]
AIAI
University of Edinburgh

Mandy Haggith
Dept. of Artificial Intelligence
University of Edinburgh

AIAI-TR-193

October 1993

This article was published in the journal "Expert Systems with Applications", 9,2, 1995, pp. 91-101.

Artificial Intelligence Applications Institute
University of Edinburgh
80 South Bridge
Edinburgh EH1 1HN
United Kingdom



[1] Author to whom all correspondence should be addressed

**Abstract**

This paper outlines the structure and the results of a research project titled *Development of a KBS for Personal Financial Planning Guided by Pragmatic KADS* [6].[2] The aim of the project was the development of a prototype knowledge-based system for personal financial planning. The prototype system developed assists a potential private investor in selecting a professionally balanced mixture of investment products that meets all the investor's requirements with regard to the investment. As part of the specification, it was decided that the prototype is to develop a portfolio for a single client investing a one-off lump sum, assuming that the client's pension, insurance and mortage have been arranged and that no other prior investments need to be taken into account.

The development of the prototype was guided by the Pragmatic KADS methodology, a streamlined version of KADS suitable to the development of small and medium-sized systems. The paper therefore describes how Pragmatic KADS helped to guide and structure the development of the prototype. The development process showed that the models that are developed as part of KADS knowledge analysis, which view knowledge from different viewpoints, contribute both to breaking down the complexity of the problem and to ensuring a comprehensive knowledge analysis. It also showed that the top-down approach of KADS has the inherent danger of neglecting some of the low-level knowledge that may have an impact on the high-level problem-solving strategies. This paper presents a way in which this problem can be resolved by using a modified version of the spiral approach to expert system development [9], guided by KADS.

# 1 Introduction

This paper describes a research project entitled *Development of a KBS for Personal Financial Planning Guided by Pragmatic KADS*. As the name suggests, the purpose of the project was to develop a knowledge based system which assisted in the task of preparing a financial plan for an potential private investor. This requires selection of a mixture of financial products which is both professionally balanced and meets all the client's requirements with regard to the investment. This task is normally carried out by a trained sales force, and is understood by very few untrained people; personal financial planning therefore fulfils the criteria for being an "expert task". This indicates that implementing a knowledge based system to perform personal financial planning is both feasible and possibly beneficial; indeed, several knowledge based systems have previously been commercially developed to perform, or assist in performing personal financial planning. Companies who have developed such

[2]The work was carried out while the first-named author was a postgraduate student of the Department of Artificial Intelligence, University of Edinburgh.

systems include Allied Dunbar [1] and Sterling Wentworth [10] in the UK, the BHF bank in Germany [22], Applied Expert Systems Inc. [2] and Chase Lincoln First Bank [12] in the USA, and Sanwa Bank [7] and Dai-Ichi Kangyo Bank [17] in Japan. For reviews of applications in this field, see [11] or [10].

The development of the system was supported by the Pragmatic KADS methodology [13]. Pragmatic KADS is a streamlined version of the KADS methodology. KADS (**K**nowledge **A**cquisition and **D**ocument **S**tructuring) [20] is a structured approach to the development of knowledge based systems and as such is to be seen in contrast to unstructured approaches such as rapid prototyping. There have been previous attempts to perform structured analysis of the knowledge required to perform personal financial planning (e.g. the CompAssys system was developed "according to slightly modified rules of information engineering" [22]), but few records are available of such attempts, and there are no known uses of KBS-specific methods to construct an expert system for personal financial planning. This project was therefore able to achieve a methodological purpose as well as producing an implemented system: the strengths and weaknesses of the Pragmatic KADS methodology on a real-life project could be evaluated, and the modelling of the acquired knowledge could be proposed as a first step to producing useful models of approaches to financial planning, and of financial products.

In order to reduce the complexity of the problem to a tractable size for the duration of the project, the following simplifying assumptions were made:

- the investor is single;
- the investor does not have an existing portfolio that needs to be taken into account;
- the investor wants to invest a single lump sum, as opposed to a series of regular payments;
- the investor's insurance, mortgage and pension arrangements have been sorted out, and so do not need to be considered.

These assumptions restrict the number of financial products which might be relevant to the investor, but do not materially affect the process of preparing a personal financial plan.

The following sections provide a concise account of how the prototype of a knowledge-based system was developed under the guidance of Pragmatic KADS within the framework of a spiral approach to expert system development. In section 2, the underlying ideas and principles of both the KADS and the Pragmatic KADS methodology are outlined. In section 3, the most important results of the knowledge acquisition and elicitation stage are described. In section 4, the knowledge and understanding of the task obtained in the previous sections is analysed with the aid of KADS models. In section 5, the most important part of the design stage of systems guided by Pragmatic KADS, the stage of behavioural design, is discussed in some detail. Section 6 explains why an iteration of knowledge acquisition, analysis and design became necessary when implementation of the system was started. Finally, in section 7, the initial aspirations are contrasted with the final achievements of the project and the usefulness of the Pragmatic KADS methodology for the development of a knowledge-based system for personal financial planning is assessed.

# 2 KADS and Pragmatic KADS

According to KADS, development of a knowledge based system is to be seen as a modelling process, during which models of the acquired knowledge at different levels of abstraction are developed. This is intended to provide a structured method for transforming the expertise extracted from the expert into the code of the planned knowledge-based system. During the knowledge analysis stage, the most important KADS model is the four-layer model of expertise, which divides problem-solving into four "layers":

**the domain layer:** the static concepts and relations specific to the domain;

**the inference layer:** a declarative description of the problem-solving task;

**the task layer:** a procedural description of the problem-solving task;

**the strategy layer:** the integration of different problem-solving tasks.

KADS' influence is most pervasive during the knowledge acquisition and analysis stages of expert system development. The overall KADS methodology provides considerably less direct guidance during design and implementation, except for recommending that these stages should strive to maintain the structure developed during the previous stages. Pragmatic KADS provides more support for the various design stages; for example, the stage of behavioural design is guided by "probing questions" derived from the approach of Kline and Dolins [16] which help to select the most appropriate AI implementation techniques.

# 3 Knowledge Acquisition

The first two stages in the development of a knowledge-based system, well before design and implementation can even be thought of, are knowledge acquisition and analysis. Knowledge acquisition consists of the collection of the knowledge that governs the decision-making in the domain; knowledge analysis consists of the interpretation of that knowledge.

Knowledge acquisition for this project focused both on very high-level strategic decisions concerning the overall structure of the client portfolio and very low-level decisions governing the selection of particular investment products as components of a client portfolio. The knowledge used in the KBS is based on two sources:

- published literature on personal financial planning;
- a series of three knowledge elicitation interviews with a professional financial advisor.

While the written literature was useful in gaining a basic understanding of the area of personal financial planning, i.e. of the investment products available on the market and of low-level rules of thumb for deciding between different investment products, it was first and foremost the interviews with the expert that provided the necessary information about how a portfolio for a client was systematically constructed.

During the three interviews with the financial expert, several knowledge elicitation strategies, such as the 20-Questions technique [4], protocol elicitation [18], and the card sort [8] were employed. The 20-Questions technique and the protocol analysis of a case study gave some important clues about the procedural approach adopted by the expert; in particular, they showed that the expert considered a portfolio to consist of three main components, which are designed in the following sequence:

**the emergency fund:** easily accessible investments for sudden and unexpected financial crises;

**the income fund:** investments designed to provide a regular income;

**the growth fund:** investments with varying degrees of risk designed to provide capital growth.

The card sort, in contrast, focused on the low-level knowledge in the form of investment products. The card sort involved the expert sorting all important investment products – each of which was printed on a single card – according to their membership to one or more of the three above-mentioned main components of the portfolio. The result of the card sort helped to build an initial bridge between the low-level knowledge extracted from the literature and the high-level reasoning strategies employed by the expert.

The application of structured knowledge elicitation techniques such as card sorting is believed to be unique amongst KBS which perform or assist with personal financial planning. The few previous publications which mention knowledge acquisition indicate that knowledge acquisition has been performed by interviewing (e.g. [5]) and/or prototyping (e.g. [7]).

# 4 Knowledge Analysis

Whilst the focus of knowledge acquisition is on collecting and compiling knowledge from different sources by means of traditional knowledge acquisition methods, during knowledge analysis the focus shifts away from traditional knowledge acquisition methods to KADS methods as the main tools for analysing the knowledge acquired. As part of KADS-guided knowledge analysis a four-layer model of expertise is developed, including the domain, the inference, the task and the strategy layer. The development of the three layers which were used in this project are described below.

## 4.1 The Domain Layer

The domain layer consists of the static concepts and relations specific to the domain of expertise being modelled. In the domain of personal financial planning, the most important concepts are:

- the main components of a portfolio (the emergency fund, the income fund, and the growth fund);
- the investment product categories, e.g. the building-society account;
- the investment products types, e.g. the TESSA account;
- the specific investment products, e.g. the Stroud and Swindon TESSA account.

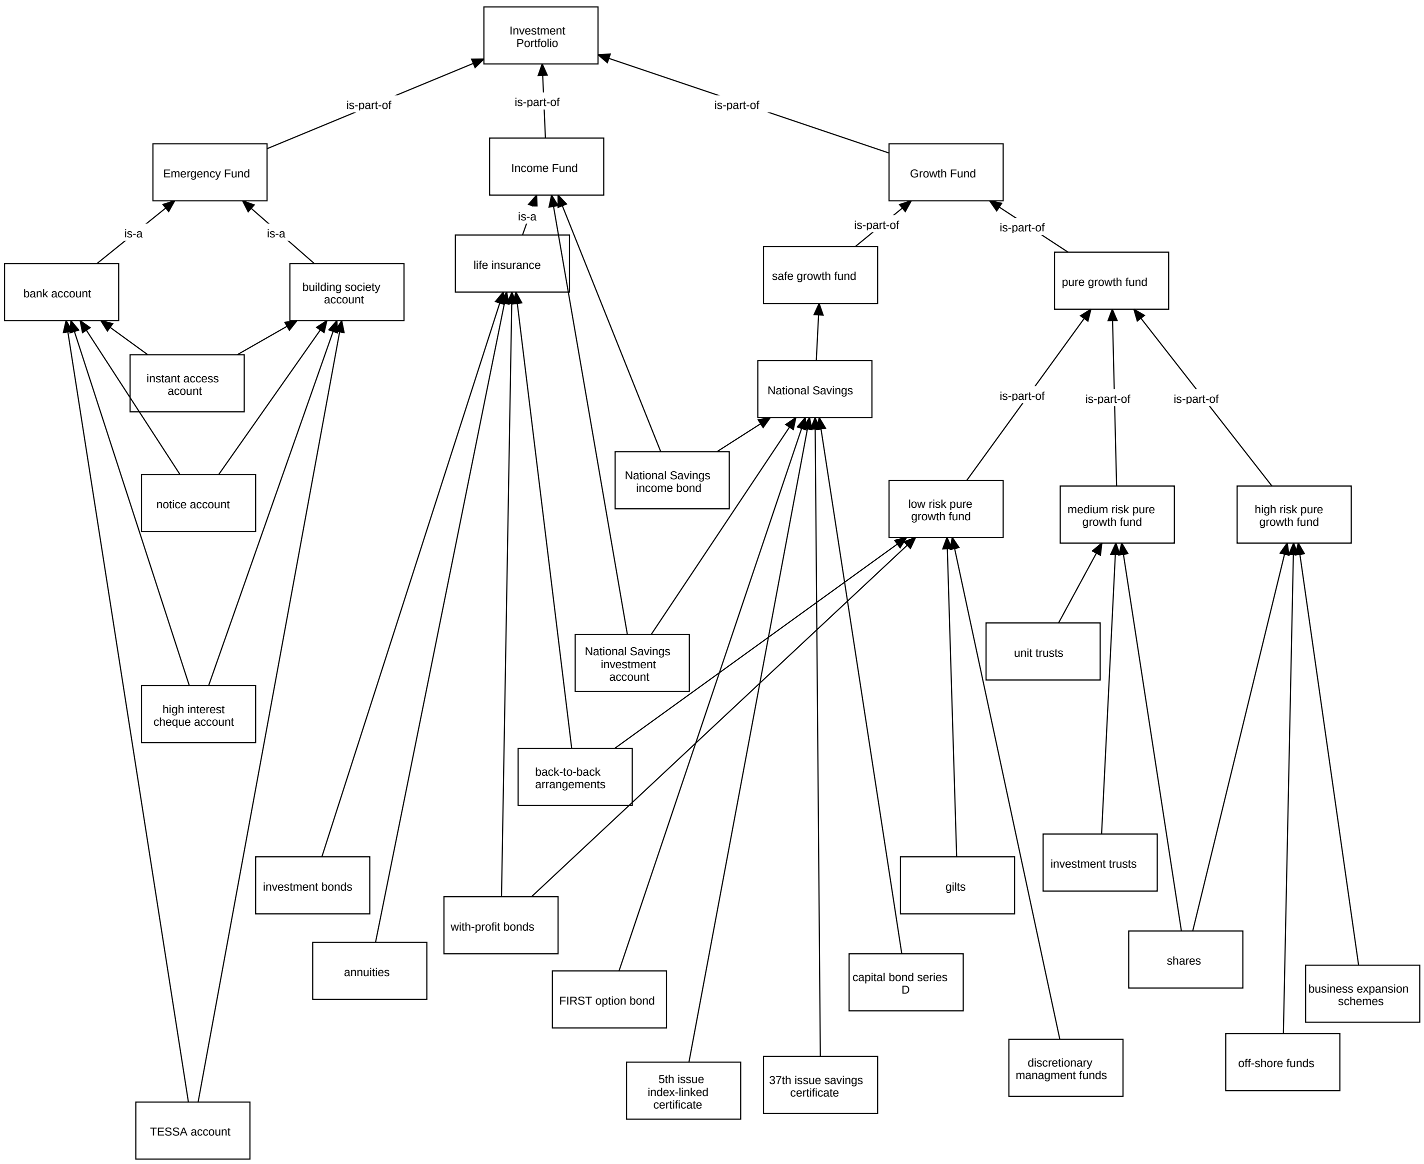


**Figure 1: Top-level Domain Model for Personal Financial Planning**

The relations between these products are either **part of**, **is a**, or **instance** relations. For example, the Stroud and Swindon TESSA account **is an instance of** a TESSA account. A TESSA account **is a** building-society account. A building-society account **is** typically used for **an** emergency fund. The emergency fund **is part of** the client portfolio.

The domain knowledge extracted from both the written literature and the interviews with the expert can be integrated into a set of “domain models”, as recommended by CommonKADS[3]. The domain of personal financial planning was modelled in a graphical manner using HARDY, a hypertext-based diagramming package [21].

The top-level domain model for the domain of personal financial planning can be seen in Figure 1. This graphical approach to domain modelling was particularly useful for three main reasons:

1. it made it possible to break down the complexity of the domain and view it at different levels of detail;

2. it was ideally suited for representing the fact that some investment products belong to more than one investment product category;

3. it allowed for easy modification and expansion of the domain model.

## 4.2 The Inference Layer

The second layer of expertise in the KADS knowledge analysis is the inference layer. Modelling of the inference layer involves two important steps: first, an *interpretation model* is selected from the KADS library of predefined interpretation models, then the selected interpretation model is adapted to the expert task at hand by instantiating it into an *inference structure*. An interpretation model is a diagrammatic declarative description of the inferences required in a typical problem-solving task[4].

The KADS library of predefined interpretation models (see figure 2) contains a number of models, each suited to a different type of task; the underlying idea is that, while the domain of each KBS tends to be very specific, experts nonetheless use domain-independent problem-solving strategies. Interpretation models fall into one of three categories: system modification, system analysis, and system synthesis. These categories are each made up of further sub-categories. For example, system synthesis comprises transformation, design, planning, and modelling. Design in turn can be further subdivided into exploration-based design (N.B. this model is unique to Pragmatic KADS – see [15]), transformation design, hierarchical design, configuration, and refinement design.

[3]CommonKADS is the successor to the KADS methodology. It is is currently under development. See [24] for details of domain modelling.

[4]Strictly speaking, there is more to an interpretation model than just the diagram of generic inference steps; however, the other components of interpretation models are used so rarely that the term "interpretation model" is often used to refer only to this diagram.

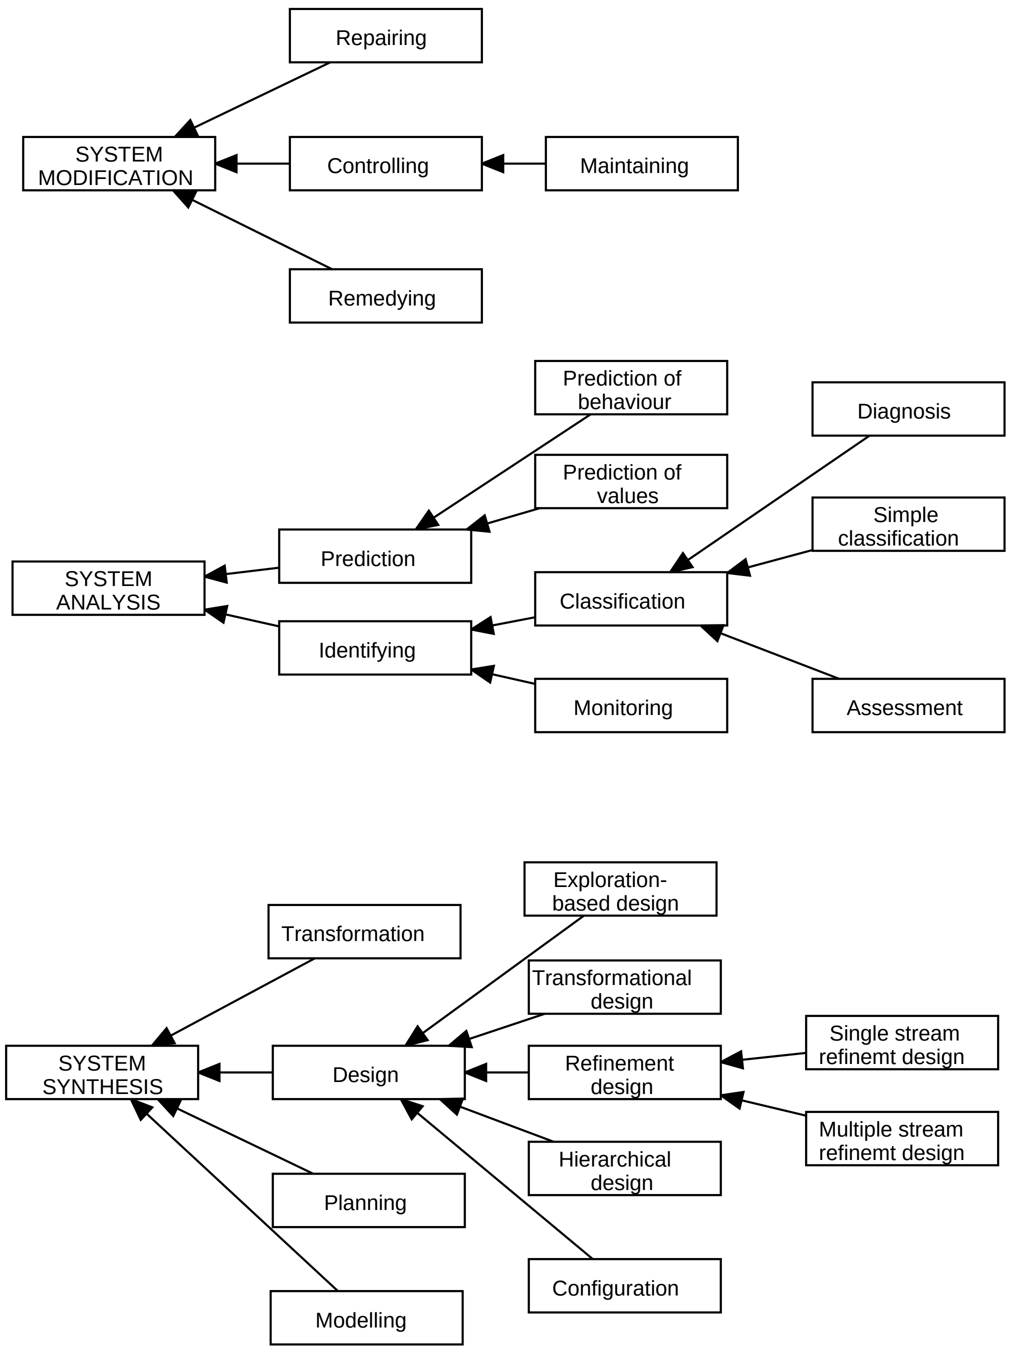


**Figure 2: Library of Interpretation models**

The idea of this hierarchically ordered library of interpretation models is that the knowledge engineer can work her way down to the most suitable interpretation model, step by step. In Pragmatic KADS, the selection of the most adequate interpretation model is also guided by a battery of questions which point to the best choice of interpretation model [13].

Based on the knowledge elicited during the interviews with the expert, the interpretation model for *hierarchical design* tasks was chosen as the most suitable model for the task of personal financial planning. This model was chosen for the following reasons:

1. The task of personal financial planning involves composing a new structural description of a possible object, an investment portfolio, within the domain; it is effectively impossible to enumerate all possible descriptions that can be developed. The task therefore falls into the category of **system synthesis**;

2. The task of personal financial planning involves designing an artefact which meets certain constraints, such as the client's requirements with regard to the purpose of the investment, but does so without a predefined framework. The task therefore belongs to the category of **design**;

3. The task of personal financial planning is tackled by breaking a financial plan into subcomponents, such as the emergency fund, the income and the growth fund, and designing each of these subcomponents indepedently. Therefore the task can be best described as one of **hierarchical design**.

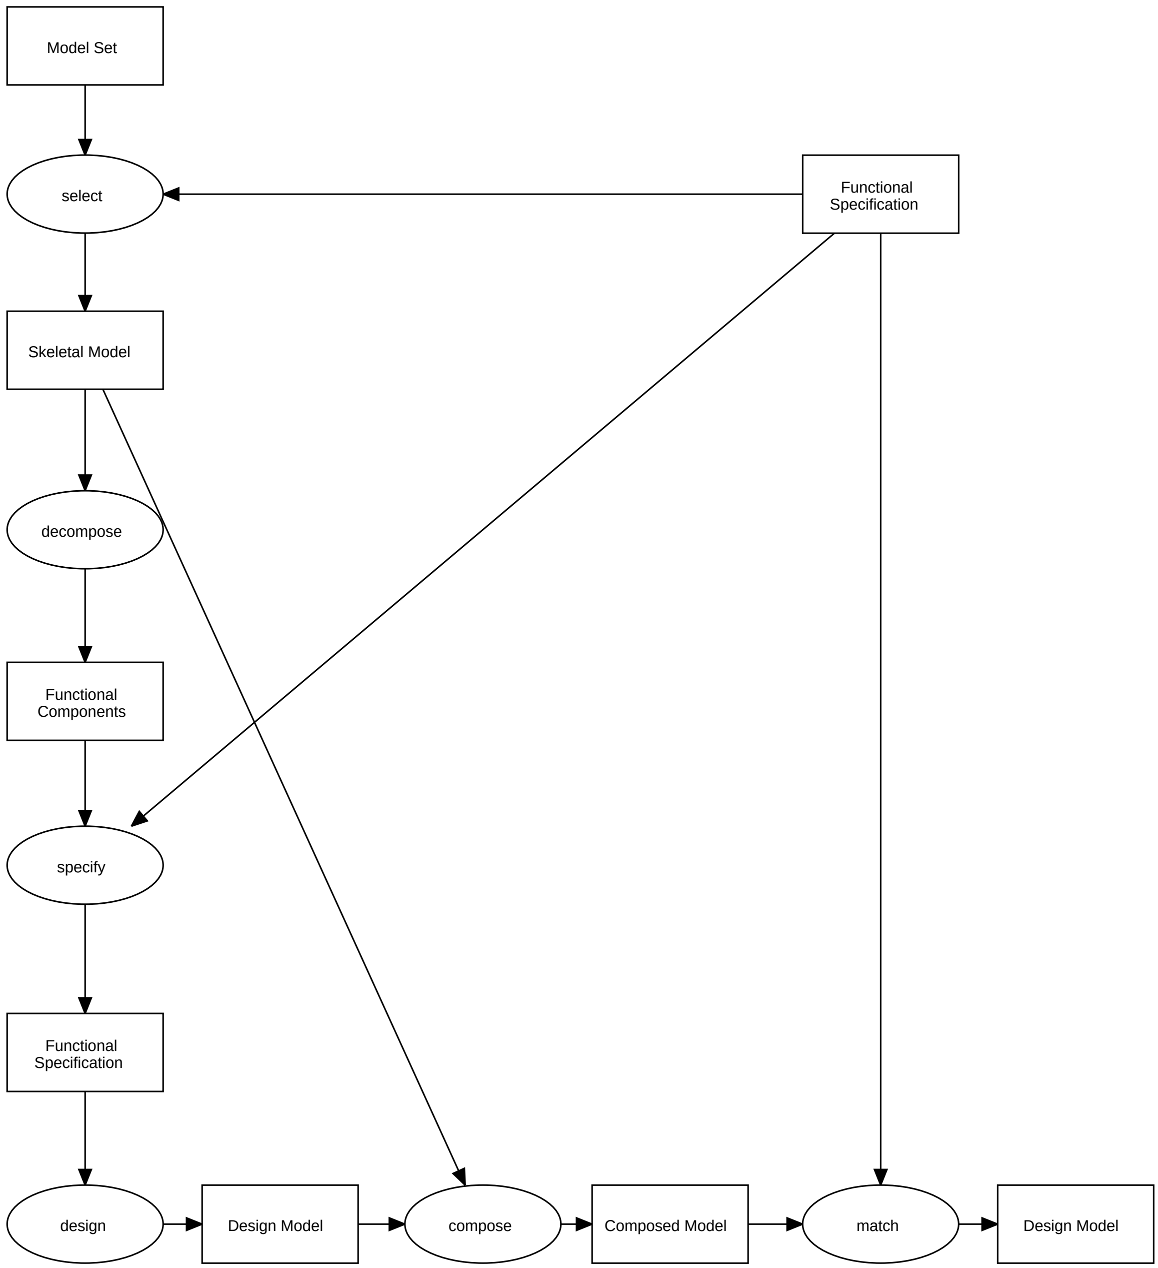


**Figure 3: The KADS Interpretation Model for Hierarchical Design**

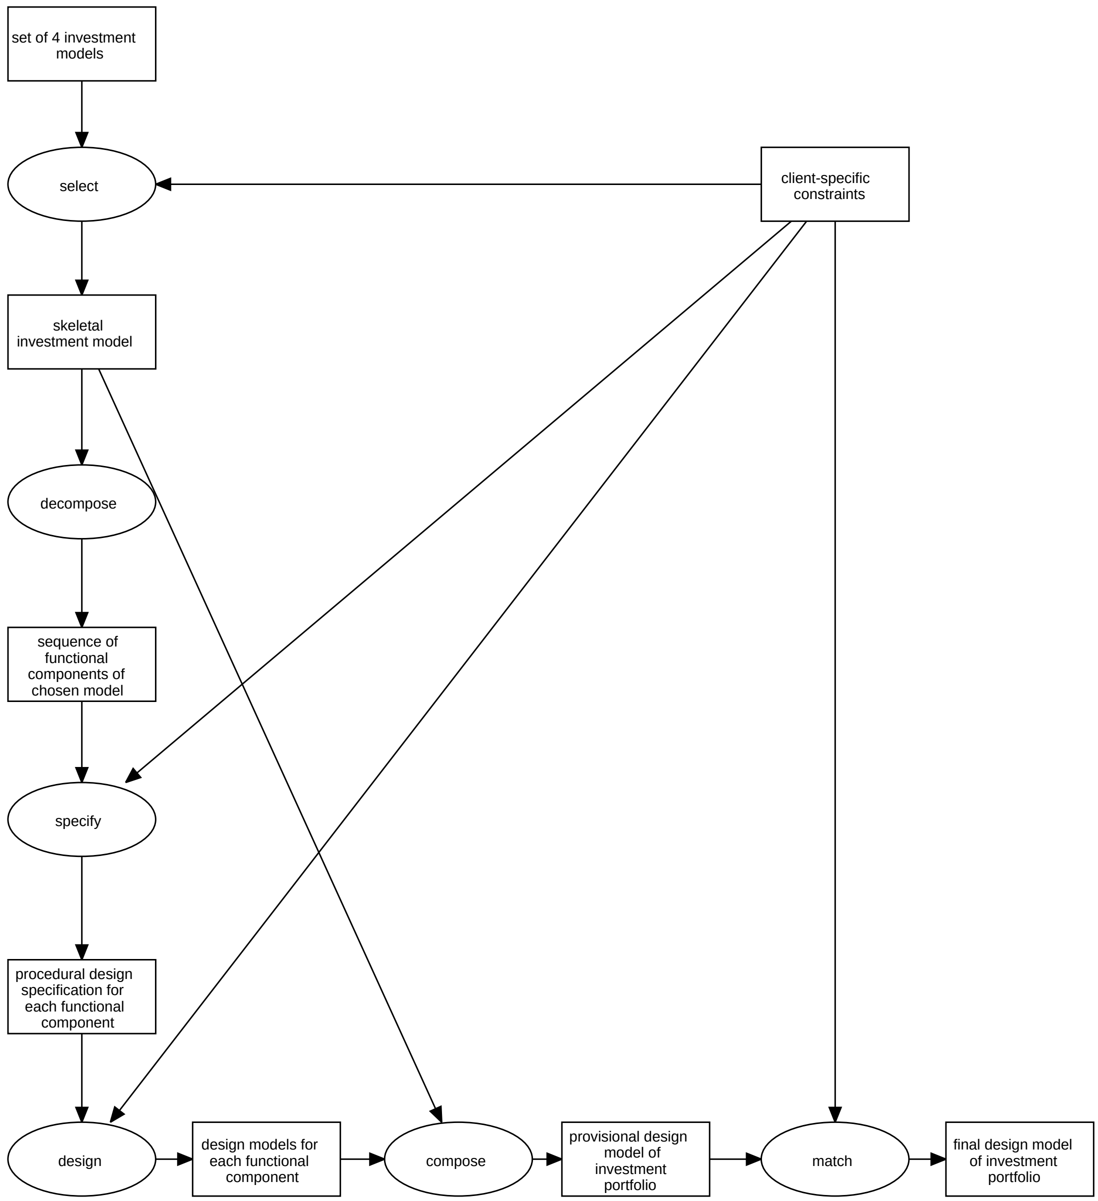


**Figure 4: The Final Inference Structure for Personal Financial Planning**

The interpretation model for hierarchical design tasks can be seen in Figure 3. This interpretation model was then instantiated into an inference structure by replacing the boxes (known as *knowledge roles*), and the ovals (known as *inference functions*) in the interpretation model with the entities of the problem domain that they represent. The resulting inference structure can be seen in Figure 4.

The inference structure in Figure 4 can be interpreted in the following manner: First of all, the client's constraints, in particular her main investment objective and the level of income she requires from the investment, are matched with the four available skeletal investment models (the income model, the mixed income/growth model, the high income mixed model, and the growth model), and the most suitable skeletal investment model is selected. This skeletal model is then decomposed into a sequence of functional components. For example, in the case of the growth model, the sequence states that an emergency fund should be allocated first, followed by a growth fund. The design procedure for each of the functional components is then specified, thereby taking into account client-specific constraints. In contrast to the predefined interpretation model, the client-specific constraints are also considered in the next step, when each functional component is designed. After all functional components have been designed, they are composed to form a provisional investment portfolio, which, after being matched with further client-specific constraints, becomes the final tailor-made investment portfolio.

## 4.3 The Task Layer

Although the interpretation model shows the inference steps performed during a particular task, it does not give any indication of the exact procedural ordering of the inferences. Such procedural ordering is highlighed by the task model which provides a procedural description of the task and as such offers a different viewpoint of the task.

On the basis of the information provided by the expert in the initial interview, the following task model was developed:

**Task** design-investment-portfolio-for-client
**Goal** assist a private investor in designing a portfolio tailored to his requirements

> **Begin** by obtaining initial information from the client
> **Select** the skeletal investment model that best meets the client's requirements from the set of investment models
> **Decompose** the chosen investment model into a sequence of functional components
> **Specify** each component's design procedure
> > For each functional component
> > > **Design** each functional component according to the component design procedure, hereby paying attention to client-specific requirements
>
> **Compose** and link the functional components to form a provisional version of the investment portfolio, hereby paying attention to the skeletal

model
**Match** the provisional version of the investment model with the client's requirements to obtain the final version of investment portfolio

**Figure 5: Task Structure for designing an investment portfolio**

# 5 Design

The software engineering life-cycle of conventional computer systems is generally characterised by the strictly sequential stages of analysis, design, and implementation. Such development is to be seen very much in contrast to the usual development of expert systems, where a prototype is developed very quickly and then repeatedly modified, expanded, validated, and tested. If KADS' aim is to introduce a more methodological approach towards the development of expert systems, then KADS should not only guide and structure the knowledge acquisition and analysis stages, but should also provide some advice on how the design and implementation of the system should proceed.

However, compared to the guidance that KADS offers during the knowledge acquisition and analysis stages, it does not give much direction on how design and implementation should progress, except for emphasizing that both stages should strive to maintain the structure of the model of expertise developed during the previous stages.

KADS-guided design therefore proceeds in a similar fashion to design for conventional computer systems and subdivides the design process into the following three stages:

1. functional design;
2. behavioural design;
3. physical design.

While functional design and physical design do not differ significantly from conventional software, it is **behavioural design** that deals with selection of knowledge representation and inference techniques, a process which differs greatly between KBS and conventional software. There is a little guidance on behavioural design in KADS; however, Pragmatic KADS offers guidance on appropriate control strategies, knowledge representations, and inference techiques. This guidance is described further below.

## 5.1 Behavioural Design

The purpose of behavioural design is to select the most suitable expert system design and AI implementation techniques for the implementation stage. In Pragmatic

KADS, this stage issupported by “probing questions”, which are partially derived from the work of Kline and Dolins [16]. This is a collection of specific questions, the answers to which help the system designer in selecting the most suitable knowledge representation and inference techniques with regard to:

- the overall architecture of the system;
- the knowledge representation strategies;
- the inference types;
- the control of the flow of inference;
- the handling of uncertainty in knowledge;
- the user interface and general input/output mechanism;
- the knowledge acquisition strategy.

In the case of the system for personal financial planning, the result of the probing questions is based both on material covered by Kline and Dolins as well as the PDQ system developed to support Pragmatic KADS [19]. The latter helps to select the most suitable AI techniques and gives a score and explanation for the method chosen to indicate how strong the supporting evidence for the relevant method is.

As a result of the probing questions, the following AI implementation techniques were selected:

**The overall architecture of the system:** Shallow reasoning, as opposed to model based reasoning, was chosen as the most suitable architectural paradigm for three reasons. First, the task can be described by a procedural representation of the expert's problem-solving steps; second, the problem space is shallow and does not require explanations in terms of cause and effect; and last, the development time is less than six months.

**The knowledge representation strategies:** Both rules and frames were chosen as knowledge representation techniques. Rules are used to represent the large and poorly structured collection of independent heuristics that can be found in the domain of personal financial planning. Most of these heuristics can be described in terms of if-then statements and are therefore particularly suitable for a representation in form of rules. Frames are used to represent the well-structured, hierarchically ordered set of static knowledge of domain concepts that was modelled in the domain layer. The format of these hierarchically ordered facts can be described in terms of slots for minimum/maximum investment, interest paid, tax liability, and similar factors and is thus suitable for representation by means of an inheritance technique, i.e. frames.

**The inference type:** Data-driven reasoning, as opposed to goal-directed reasoning, was chosen as the direction in which inference should proceed. Data-driven reasoning is suitable because a lot of facts, such as client details, are available at the start. These facts are then used to find a single solution amongst the great number of solutions that might be possible. Because of both the availability of data and the possibly very large number of solutions, data-driven reasoning was selected as the most suitable inference type. In rule-based KBS, data-driven reasoning is typically implemented using forward chaining rules.

**The control of flow of inference:** Regarding the control of the flow of inference, it was decided that search should proceed in a depth-first rather than a breadth-first manner. The main reason for this decision is that the search tree for product selection (see Fig. 1) has a high branching factor and consists of only four levels of depth – the fund type, the investment product category, the investment product type, and the investment product (the latter not shown in the tree). Due to the high branching factor and the low number of levels in the tree, depth-first search is more suitable than breadth-first search.

**The handling of uncertainy:** In general, there is little uncertainty in the task of personal financial planning. The exception to the rule is the development of the high income mixed model, where it may become necessary to decompose the portfolio because it does not achieve the income needed. Since this involves falsifying assumptions, a truth maintenance system may be worth considering. However, this was not included in the prototype.

**The user interface and general input/output mechanism:** Regarding the user interface, a multiple-selection menu was seen as an adequate way of eliciting questions from the user. In a way this corresponds to the questionnaire that the financial advisor gives to his client to complete. Furthermore, one should consider the possibility of a user-client partnership, where the user participates in the choice of investment products.

**Knowledge acquisition:** If the system was to include automatic knowledge acquisition, which the prototype does not, the best way to do this would be to use rule induction or case-based reasoning. Rule induction is appropriate because the decisive factors in most financial problem cases can be easily identified. Case-based reasoning is appropriate, because there is a vast collection of portfolios that have been developed over the years and it should be easily possible to extract the most influential features from these.

On the basis of the results of the probing questions, the expert system shell CLIPS 5.1 was chosen as the most suitable tool for implementation, because CLIPS allows knowledge to be represented as rules and frames, CLIPS' rules are designed for forward-chaining, and CLIPS was readily available.

## 6 The Necessity to Iterate Knowledge Acquisition, Analysis and Design

Conventional system development normally expects software development to follow a linear process, where the analysis stage is completed before design is started, and the design stage is completed before implementation stage. KADS has a similar emphasis, with design being heavily based on the results of the knowledge analysis, and implementation being dependent on the existence of a KBS design.

However, during the implementation of this prototype expert system for personal financial planning, several problems arose which made it necessary to return to the stages of knowledge acquisition and analysis in order to refine the models of analysed knowledge. The three main problems which arose were:

**Problems regarding the accuracy of the task model:** It turned out that the task model had been incorrect and should really read as:

> **Task** design-investment-portfolio-for-client
> **Goal** assist a private investor in designing a customised portfolio
>
> **Begin** by obtaining initial information from the client
> **Select** from the investment model set the model that best meets the client's requirements
> **Decompose** the chosen skeletal investment model into a sequence of functional components
> **Specify** each components design procedure
>   For each functional component
>     **While** current functional component is not fully designed
>       **Specify** one part of the relevant functional component's design procedure
>       **Design** the part specified, hereby paying special attention to client requirements
> **Compose** and link the functional components to form a provisional version of the investment portfolio, hereby paying attention to the skeletal model
> **Match** the provisional version of the investment model with the client's requirements to obtain the final version of investment portfolio

**Figure 6: Revised Task Structure for designing an investment portfolio**

In other words, the task model was changed to reflect the fact that the expert specifies and designs one sub-component before she goes on to specify and design the next sub-component.

**Problems regarding the best product choice:** The products in the investment category life insurance are all characterised by leaving it largely up to the client how much income he wants to withdraw from them. This high degree of flexibility made it difficult to differentiate between the different products. A new slot, *purpose*, was needed to make an informed choice when choosing between specific products.

**Problems regarding the set of investment models:** The emergence of increasing similarities between the three income-related models (the income model, the mixed model, and the high income mixed model) called for a review of the validity of each model. In the final analysis, however, it was decided that all three models were needed, as the motivation underlying each model was fundamentally different and made use of different heuristics.

The best way to describe this iteration is to view it as a KADS-guided spiral approach to KBS development. This can be illustrated using a slightly modified version of the spiral model of expert system development proposed by [9] (see Figure 7).[5]

As Figure 7 shows, the modified spiral model assumes that one pass through the stages of knowledge acquisition, knowledge analysis, design and implementation may not be sufficient. Rather, a second, and possibly more passes may be required.

The necessity for this loop in system development arose for the following two main reasons:

1. In contrast to "rapid prototyping" methods for KBS development which are empirical and inherently work bottom-up, KADS takes a top-down approach based on generic interpretation models. While a top-down approach has many advantages for guidance and structuring, it has the inherent danger of neglecting and/or overlooking important low-level knowledge on the first pass, and of failing to identify how low-level knowledge interacts with higher-level decisions.[6]

[5] The successor to the KADS methodology, known as CommonKADS, recommends moving away from a linear approach to KBS development towards a spiral approach to KBS project management – see [3]. The results of this paper provide one of the first empirical examples of the need for this approach.

[6] Again, CommonKADS is addressing this area, by introducing an abstract domain model which explicitly links the domain and inference levels of the model of expertise. See [14] for a worked example.

2. The domain knowledge was more complex than expected. In particular, the declarative knowledge in the domain of personal financial planning turned out to be less structured than expected, and the procedural knowledge in the domain of financial planning turned out to be more complex and sophisticated than expected. The task of design was also made more complex by by the variability in both the number and size of the components of a typical portfolio.

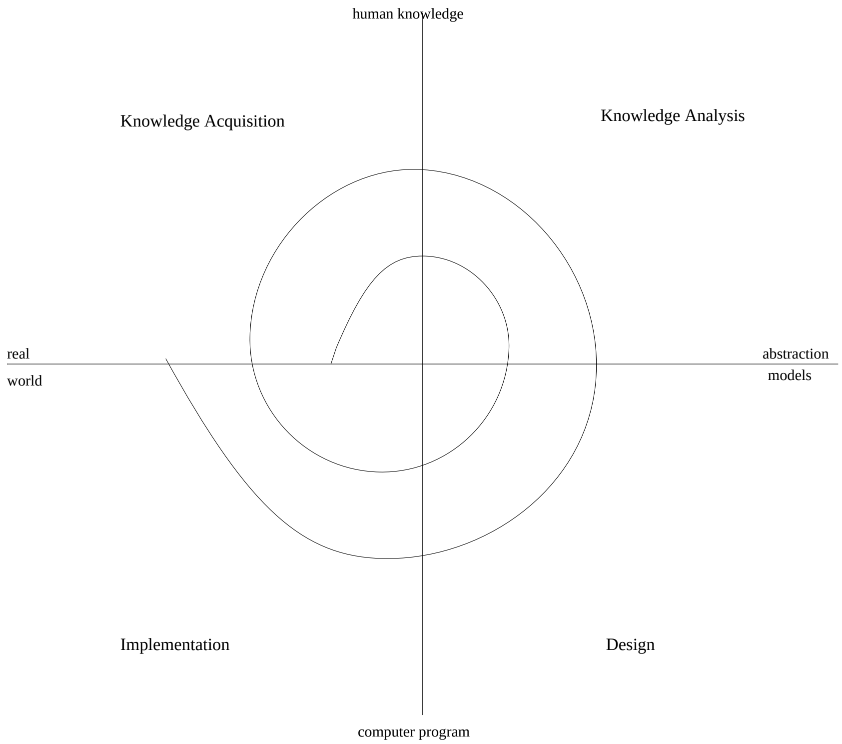


**Figure 7: KADS-guided Spiral Model of Expert System Development**

In the second pass through the spiral development process, the problems that arose during implementation were resolved, and system development then smoothly proceeded with design and the implementation of the prototype system. Similarly, the progression of the prototype system into a fully-fledged knowledge-based system, and the maintenance of the highly changeable knowledge base of the system could well be guided by further passes through the stages of knowledge acquisition, analysis, design, and implementation, guided by the spiral model.

Therefore it is worth noting that while this loop in system development was unexpected, it does not negate the validity and usefulness of KADS. On the contrary, the use of KADS models helped to highlight those areas where iteration is necessary.

# 7 Conclusions

As mentioned in the introduction, the initial aim of the project was to develop a prototype knowledge-based system for personal financial planning which would assist a potential private investor in selecting a professionally balanced mixture of investment products that would meet all the investor's requirements with regard to the investment. The development of this prototype was to be guided by the Pragmatic KADS methodology.

On the practical side, the prototype expert system for financial planning was implemented. Due to time restrictions only the three income-related models were implemented, whilst the growth model was not implemented. The prototype system as it stands is fully functional, but could be extended in several ways, in particular by

- the incorporation of a growth portfolio;
- a more comprehensive treatment of the client's financial situation prior and subsequent to the investment;
- extensions regarding the static domain knowledge in form of specific investment products.

On the theoretical side, the development of the prototype expert system was to be guided by the use of Pragmatic KADS. It has been shown that the contribution of Pragmatic KADS was most influential and pervasive during the knowledge acquisition and analysis stages. During knowledge acquisition, the KADS models served as a useful communication tool in the interviews with the expert. During knowledge analysis, the model of expertise laid the foundations for the structure of the system to be developed and had a decisive, though indirect, impact on the design and implementation stages. During behavioural design, the use of Kline and Dolins' *probing questions* contributed greatly to selecting the most appropriate AI implementation techniques.

Although the development of the system did not proceed in the linear way initially expected, the KADS models helped to highlight those areas where more knowledge acquisition and knowledge analysis was needed to arrive at a complete, consistent and correct implementation.

Pragmatic KADS is therefore considered to be a valuable tool in the hands of KBS developers. Its greatest strengths lie in the provision of a structure which can be used as a starting point for system development; in the separation of different types of knowledge into different models; and in the use of these models for communication with an expert. On the other side of the coin, KADS leans towards a sequential model of KBS analysis, design and implementation, which was inappropriate for this project; and guidance on the design stage could usefully be developed further.

It should be noted that CommonKADS, the successor to KADS, was still under development at the time of writing, and is expected to produce considerable guidance on different approaches to model development, as well as recommending refinements to some of the models. See [23] for details.

# References

Andrews, B. (June 1989). Successful Expert Systems. In *Financial Times Business Information*, pages 49–50.

Anon. (June 1987). APEX changes tack. *Expert Systems User*, 2,11:5.

Bright, C., Martil, R., Williams, D. and Rajan, T. (1991). The KADS-II Framework for KBS Project Management. ESPRIT Project P5248 KADS-II, Lloyds Register.

Burton, A.M., Shadbolt, N.R., Rugg, G. and Hedgecock, A.P. (1988). Knowledge Elicitation Techniques in Classification Domains. In *Proceedings of ECAI-88: The 8th European Conference on Artificial Intelligence.*

Decio, E., Spampinato, L. and di Pasquale, S. (June 1988). PECUNIA: an expert portfolio advisor. In *Proceedings of the 8th international conference on Expert Systems and their Applications*, pages 451–461, Avignon.

Dirks, S. (Sept 1993). Development of a knowledge-based system for personal financial planning. Unpublished M.Sc. thesis, Dept of Artificial Intelligence, University of Edinburgh.

Feigenbaum, E., McCorduck, P. and Nii, P. (1988). *The Rise of the Expert Company.* Macmillan.

Gammack, J.G. (1987). Different techniques and different aspects on declarative knowledge. In Kidd, A.L., (ed.), *Knowledge Acquisition for Expert Systems.* Plenum Press.

Giarratano, J. and Riley, G. (1989). *Expert Systems - Principles and Programming.* PWS-Kent Publishing Company.

Guilfoyle, C. and Jeffcoate, J. (February 1988). In *Expert Systems in Banking & Securities*, page 156. Ovum.

Humpert, B. and Holley, P. (May 1988). Expert Systems in finance planning. *Expert Systems*, 5,2:79.

Kindle, K., Cann, R., Craig, M. and Martin, T. (1989). PFPS: Personal Financial Planning System. In *Innovative Applications of Artificial Intelligence*, pages 51–61. AAAI.

Kingston, J.K.C. (1993a). Pragmatic KADS 1.0. Technical Report AIAI-IR-13, AIAI, University of Edinburgh.

Kingston, J.K.C. (1993b). Re-engineering IMPRESS and X-MATE using CommonKADS. In *Expert Systems 93.* British Computer Society, Cambridge University Press.

Kingston, J.K.C. (1994). Design by Exploration: A Proposed CommonKADS Inference Structure. *Submitted to 'Knowledge Acquisition'*.

Kline, P.J. and Dolins, S.B. (1989). *Designing expert systems : a guide to selecting implementation techniques*. Wiley.

Komahashi, K. (1989). Japanese wisdom gets electric. *Best of Business International*, pages 32–37.

Kuipers, B. and Kassirer, J.P. (1987). Knowledge acquisition by analysis of verbatim protocols. In Kidd, A.L., (ed.), *Knowledge Acquisition for Expert Systems: A Practical Handbook*, chapter 3, pages 45–71. Plenum Press.

MacNee, C. (Sept 1992). PDQ: A knowledge-based system to help knowledge-based system designers to select knowledge representation and inference techniques. Unpublished M.Sc. thesis, Dept of Artificial Intelligence, University of Edinburgh.

Schreiber, A. Th., Wielinga, B. J. and Breuker, J. A., (eds.). (1993). *KADS: A Principled Approach to Knowledge-Based System Development*. Academic Press, London.

Smart, J. (April 1993). HARDY. *Airing*, 15:3–7. AIAI, University of Edinburgh.

Thomczyk, C. (March 1991). Knowledge-based Asset Structure Analysis. In *An Analysis of Expert Systems Applications in the European Banking Industry*. Universita Cattolica, Milan.

Wielinga, R. (October 1993). Expertise Model: Model Definition Document. CommonKADS Project Report, University of Amsterdam, KADS-II/M2/UvA/026/2.0.

Wielinga, R., van de Velde, W., Schreiber, G. and Akkermans, H. (Jun 1993). Expertise Model Definition Document. CommonKADS Project Report, University of Amsterdam.

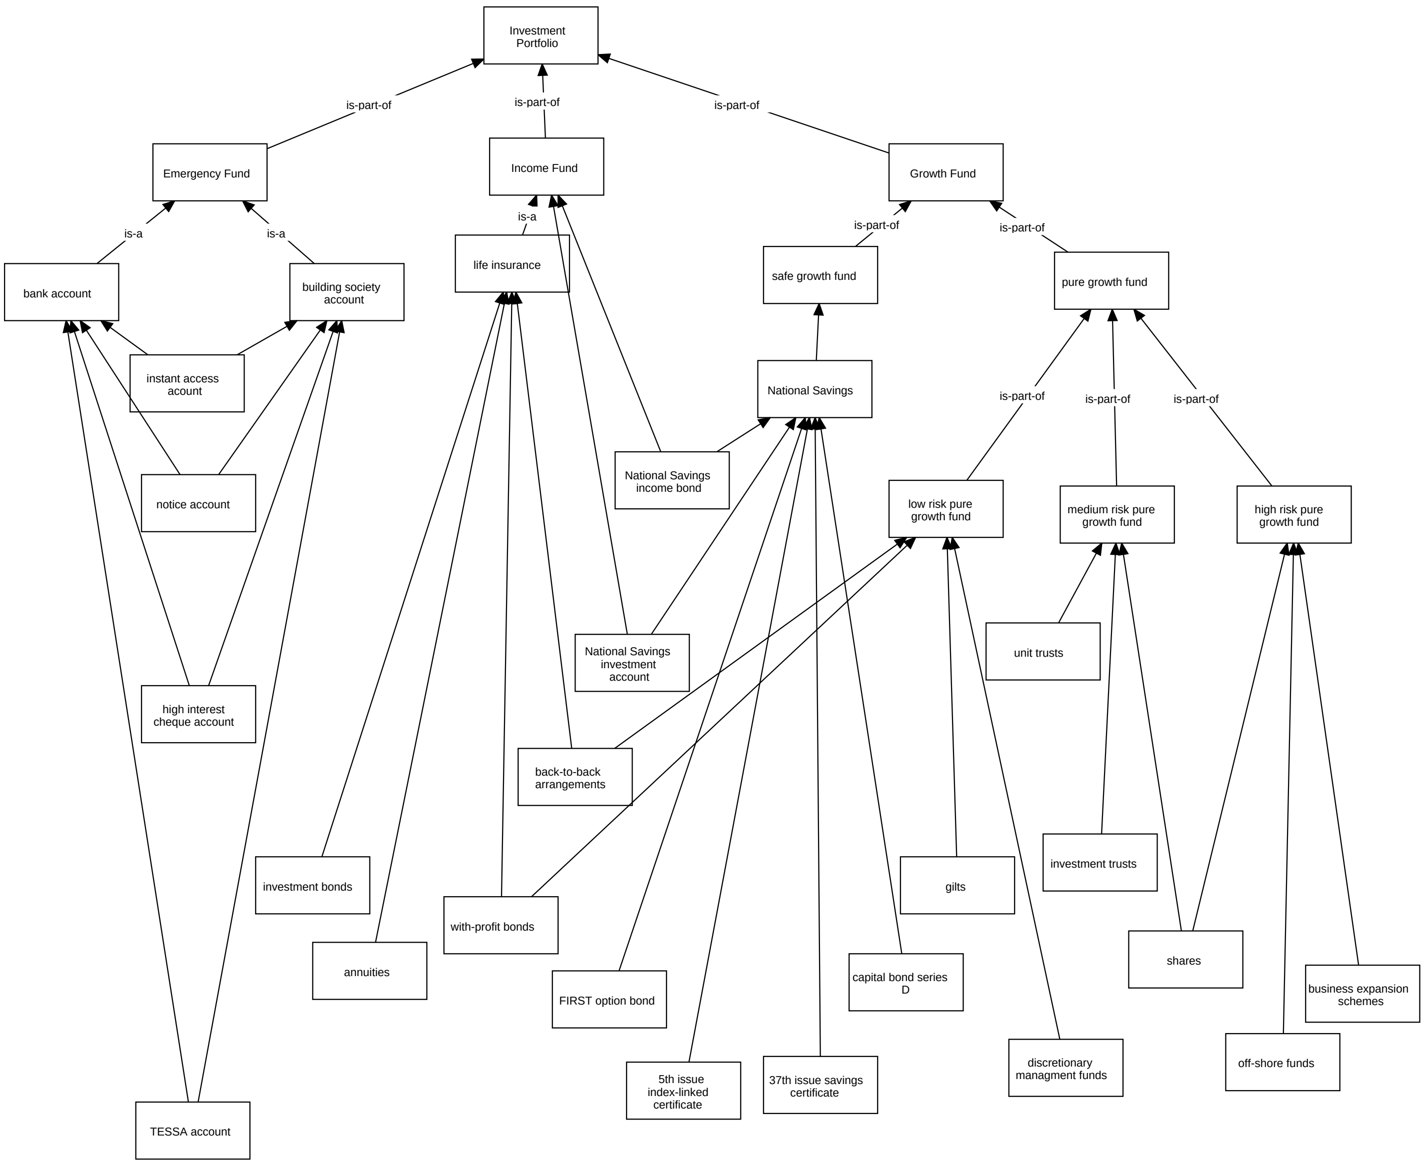

Investment Portfolio
is-part-of
is-part-of
is-part-of
Emergency Fund
Income Fund
Growth Fund
is-a
is-a
is-a
is-part-of
is-part-of
bank account
building society account
life insurance
safe growth fund
pure growth fund
instant access acount
National Savings
is-part-of
is-part-of
is-part-of
notice account
National Savings income bond
low risk pure growth fund
medium risk pure growth fund
high risk pure growth fund
National Savings investment account
unit trusts
high interest cheque account
back-to-back arrangements
investment trusts
investment bonds
gilts
with-profit bonds
shares
annuities
capital bond series D
business expansion schemes
FIRST option bond
discretionary managment funds
off-shore funds
5th issue index-linked certificate
37th issue savings certificate
TESSA account

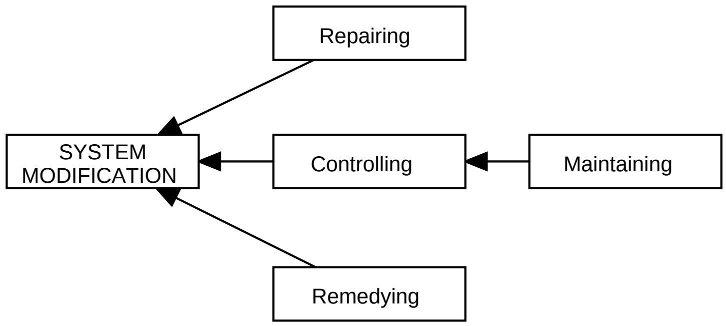
Repairing
SYSTEM MODIFICATION
Controlling
Maintaining
Remedying

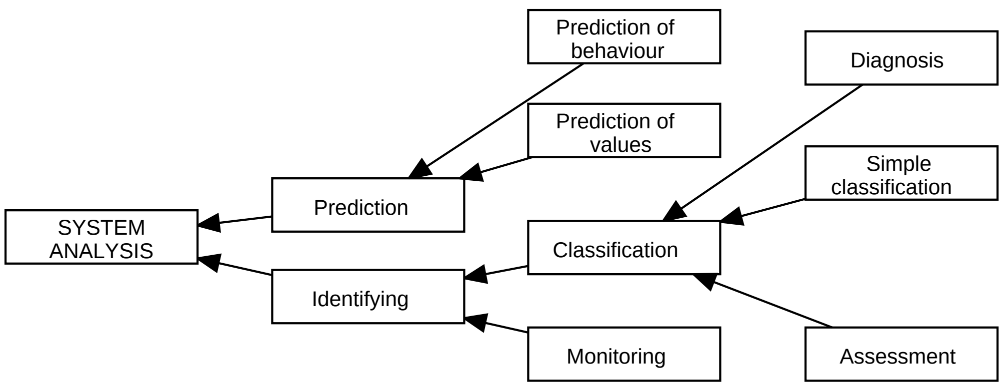
Prediction of behaviour
Diagnosis
Prediction of values
Simple classification
Prediction
SYSTEM ANALYSIS
Classification
Identifying
Monitoring
Assessment

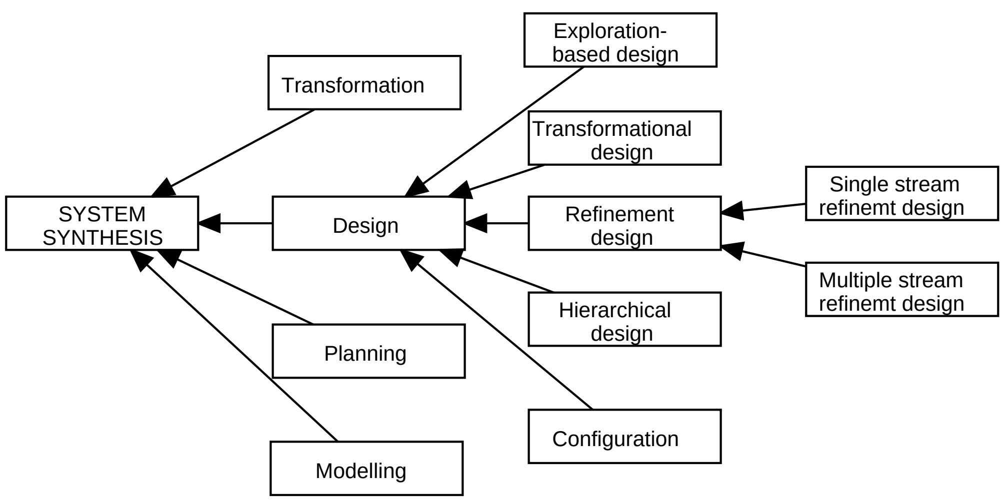
Exploration-based design
Transformation
Transformational design
Single stream refinemt design
SYSTEM SYNTHESIS
Design
Refinement design
Multiple stream refinemt design
Hierarchical design
Planning
Configuration
Modelling

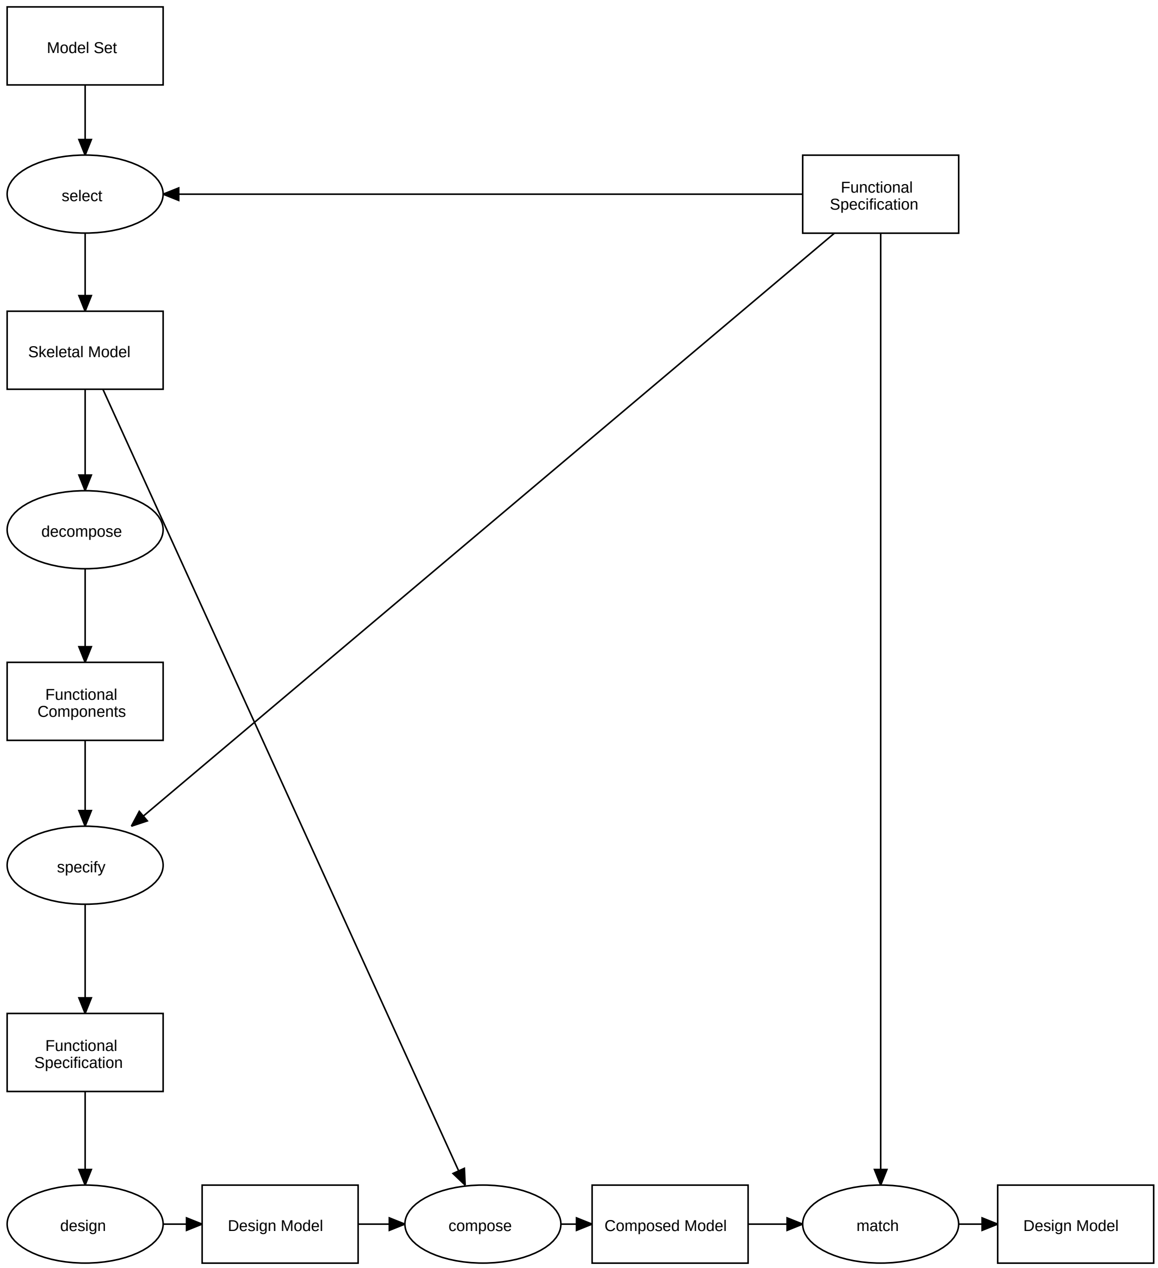
Model Set
select
Functional
Specification
Skeletal Model
decompose
Functional
Components
specify
Functional
Specification
design
Design Model
compose
Composed Model
match
Design Model

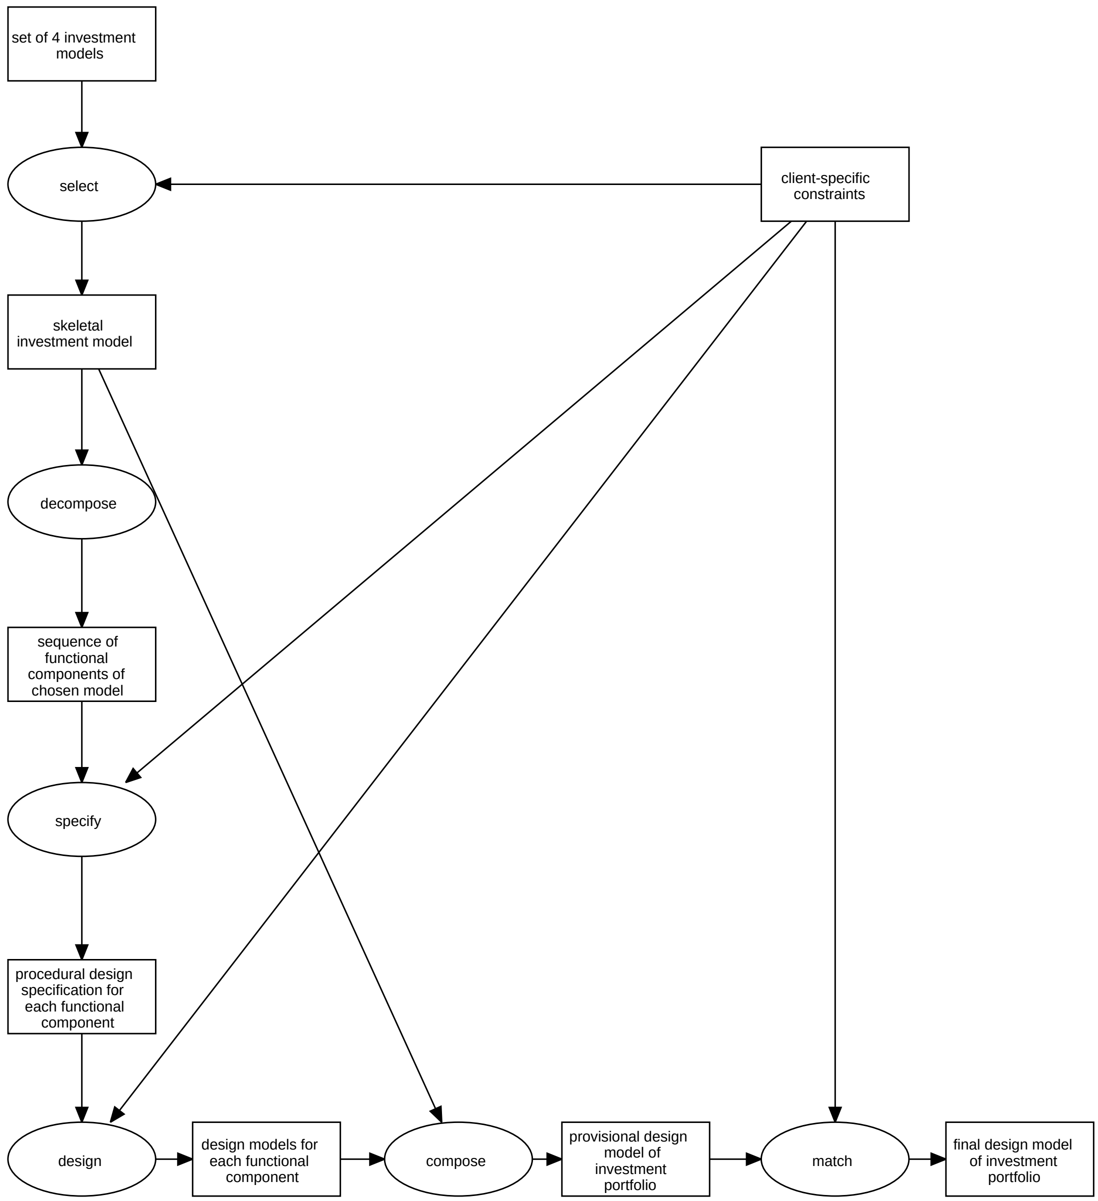
set of 4 investment models
select
client-specific constraints
skeletal investment model
decompose
sequence of functional components of chosen model
specify
procedural design specification for each functional component
design
design models for each functional component
compose
provisional design model of investment portfolio
match
final design model of investment portfolio

## References


[1] B. Andrews. Successful Expert Systems. In *Financial Times Business Information*, pages 49–50, June 1989.

[2] Anon. APEX changes tack. *Expert Systems User*, 2,11:5, June 1987.

[3] C. Bright, R. Martil, D. Williams, and T. Rajan. The KADS-II Framework for KBS Project Management. ESPRIT Project P5248 KADS-II, Lloyds Register, 1991.

[4] A.M. Burton, N.R. Shadbolt, G. Rugg, and A.P. Hedgecock. Knowledge Elicitation Techniques in Classification Domains. In *Proceedings of ECAI-88: The 8th European Conference on Artificial Intelligence*, 1988.

[5] E. Decio, L. Spampinato, and S. di Pasquale. PECUNIA: an expert portfolio advisor. In *Proceedings of the 8th international conference on Expert Systems and their Applications*, pages 451–461, Avignon, June 1988.

[6] S. Dirks. Development of a knowledge-based system for personal financial planning. Master's thesis, Dept of Artificial Intelligence, University of Edinburgh, Sept 1993.

[7] E. Feigenbaum, P. McCorduck, and P. Nii. *The Rise of the Expert Company.* Macmillan, 1988.

[8] J.G. Gammack. Different techniques and different aspects on declarative knowledge. In A.L. Kidd, editor, *Knowledge Acquisition for Expert Systems.* Plenum Press, 1987.

[9] J. Giarratano and G. Riley. *Expert Systems - Principles and Programming.* PWS-Kent Publishing Company, 1989. A good introductory text with a practical approach, which also supplies a disk of the CLIPS programming tool, and a manual on the use of CLIPS.

[10] C. Guilfoyle and J. Jeffcoate. In *Expert Systems in Banking & Securities*, page 156. Ovum, February 1988.

[11] B. Humpert and P. Holley. Expert Systems in finance planning. *Expert Systems*, 5,2:79, May 1988.

[12] K. Kindle, R. Cann, M. Craig, and T. Martin. PFPS: Personal Financial Planning System. In *Innovative Applications of Artificial Intelligence*, pages 51–61. AAAI, 1989.

[13] J.K.C. Kingston. Pragmatic KADS 1.0. Technical Report AIAI-IR-13, AIAI, University of Edinburgh, 1993.

[14] J.K.C. Kingston. Re-engineering IMPRESS and X-MATE using CommonKADS. In *Research and Development in Expert Systems X*, pages 17–42. Cambridge University Press, 1993. http://www.aiai.ed.ac.uk/ jkk/publications.html.

[15] J.K.C. Kingston. Design by Exploration: A Proposed CommonKADS Inference Structure. *Submitted to 'Knowledge Acquisition'*, 1994.

[16] P. J. Kline and S. B. Dolins. *Designing expert systems : a guide to selecting implementation techniques.* Wiley, 1989.

[17] K. Komahashi. Japanese wisdom gets electric. *Best of Business International*, pages 32–37, 1989.

[18] B. Kuipers and J.P. Kassirer. Knowledge acquisition by analysis of verbatim protocols. In A.L. Kidd, editor, *Knowledge Acquisition for Expert Systems: A Practical Handbook*, chapter 3, pages 45–71. Plenum Press, 1987.

[19] C MacNee. PDQ: A knowledge-based system to help knowledge-based system designers to select knowledge representation and inference techniques. Master's thesis, Dept of Artificial Intelligence, University of Edinburgh, September 1992.

[20] A. Th. Schreiber, B. J. Wielinga, and J. A. Breuker, editors. *KADS: A Principled Approach to Knowledge-Based System Development.* Academic Press, London, 1993.

[21] J. Smart. HARDY. *Airing*, 15:3–7, April 1993. AIAI, University of Edinburgh.

[22] C. Thomczyk. Knowledge-based Asset Structure Analysis. In *An Analysis of Expert Systems Applications in the European Banking Industry.* Universita Cattolica, Milan, March 1991.

[23] B. Wielinga. Expertise Model: Model Definition Document. CommonKADS Project Report, University of Amsterdam, October 1993. KADS-II/M2/UvA/026/2.0.

[24] B. Wielinga, W. van de Velde, G. Schreiber, and H. Akkermans. Expertise Model Definition Document. CommonKADS Project Report, University of Amsterdam, Jun 1993.