Common Sense Investing: Bridging the Gap Between Expert and Novice

Ashwani Kumar
MIT Media Laboratory
Cambridge, MA
ashwani@mit.edu
617-253-3716

Sharad C. Sundararajan
IBM, Design Automation
Fishkill, New York
sharads@us.ibm.com
781-899-0341

Henry Lieberman
MIT Media Laboratory
Cambridge, MA
lieber@media.mit.edu
617-253-0315

ABSTRACT

In this paper, we describe Common Sense Investing (CSI), an interactive investment tool that uses a knowledge base of common sense statements in conjunction with domain knowledge to assist personal investors with their financial decisions, primarily asset-allocation. In interfaces that provide expert advice, one key problem is *elicitation* – how to ask questions that enable the expert model to make decisions, and at the same time, are understandable to the novice. The second problem is *explanation* – how to explain rationale behind expert decisions in terms that the user can understand. Many programs already encode expert models, but few have good models of novice knowledge, especially where broad knowledge of everyday life might bear on the subject. OMCSNet [1], a semantic network representation of the OpenMind Common Sense Knowledge Base [2], is the source of a wide range of facts about day-to-day life. CSI maps the user's goals, expressed in concepts from OMCSNet, to the expert's goals, expressed in technical financial terms. Instead of asking "What's your tolerance for risk?", where the user might not understand the concept of risk tolerance, we can ask "Do you currently have a lot of credit card debt?". Aligning the expert's questions and decisions with common sense knowledge pertinent to the user increases the user's confidence in the ability of the system to meet their needs.

Keywords

CSI, Common Sense Reasoning, investment, frames, information filter

ACM Classification Keywords

H.5.2. User Interfaces. H.5.3. Web-based interaction

INTRODUCTION

There exist numerous investment tools [3,4,5,6,7,8,9] that claim to come up with the best strategies for asset allocation through a sequence of questions to gauge the user's inclination towards investment and willingness to take risks. But in our research what we found was the lack of sufficient control to the user, limited personalization and limited usability scenarios. The tools need a richer interactive experience. These tools do not seem to exploit commonsense knowledge to either achieve the user's goal or to make the interaction more natural. Common sense, as we understand it in today's world is shared knowledge that puts everyone on the same page and provides an enormous closeness to human thought, hence making communication easier and more intuitive. Common sense can be used to explain the success and failure of different scenarios and help troubleshoot problems along the lines of Woodstein [10], an interactive debugger. There is a huge effort going on in this direction at the MIT Media Lab, where they are building a database with millions of common-sense facts that people come across in everyday life [1,2].

In this project there is an attempt to bridge the gap between naïve and expert knowledge systems by providing an intuitive interactive framework where the user can interact with the system using natural language sentences without being overwhelmed by the expert knowledge processing that the system performs. For instance, lay users cannot objectively specify their risk tolerance, as they may not be aware of the repercussions of taking low or high risk, so it is essential for the tool to engage in a dialogue and gauge the users' risk tolerance. In the following section, we describe the functional architecture followed by brief descriptions of various underlying components in subsequent sections.

ARCHITECTURE

The system architecture consists the following components:

1. Common Sense Analyzer (CSA)
2. Expert financial engine
3. InfoFilter -Information Filter
4. OMCS Interface
5. Investment Strategies

The investment strategist interprets the analyzed request from the CSA and accumulates relevant information from the InfoFilter and Expert Financial Engine. The CSA plays a crucial role in bridging the gap between natural man-machine interactions and expert system processing. CSA comprises of a Natural Language Understanding front-end, which processes the user's commands in natural language to extract investment and goal semantics. This abstract level semantics is correlated with the common sense knowledge base in order to establish various goal and action dependencies. The CSA implements an interface, which interacts with the OMCSNET using SOAP (Simple Object Access Protocol) messages. OMCS is configured to run as a web service, which is queried to extract semantic associations in the form of predicates.

SYSTEM DESCRIPTION

The input to the system contains the investment amount, timeline and intended purpose. The output will be suggestions for asset allocations and best individual asset picks. The system contains a natural language dialogue framework, an in-built browser and a user-interface to exchange information with the system and to retrieve explanations and history of interaction, all interfaced to a common-sense database.

Common Sense Analyzer

Contemporary research in the area of interactive goal-driven systems has emphasized the importance of having a dialogue-based interaction as opposed to fixed menu or scenario based interactions with the user [17]. However, having dialogues in the mode of natural languages requires that the system have adequate language understanding capabilities, fail-soft inference and deduction mechanisms. It is imperative that the system has sufficient common sense knowledge and optimal application-specific knowledge i.e. expert knowledge. From the usability point of view, it is also desired to maintain a seamless and intuitive interface that bridges these two different types of knowledge pieces.

In CSI, our central goal is to achieve this kind of interactivity, without sacrificing application performance or overloading the end user with the application specific modus operandi. The key idea is to specify suitable mappings from natural language utterances to expert system behaviors and vice versa. The important thing to note here is that these mappings are dynamic in the sense that they evolve with interactions, they are personalized based on the user's profile, and they get refined as the common sense knowledge base gets richer.

Natural Language Understanding (NLU) Unit

All the user's requests are first tagged using a Parts of Speech (POS) Tagger. The tagged text is chunked using a text-chunker, which groups tagged words within an utterance to disjoint classes based on some pre-defined rules. Further, a semantic analyzer produces the semantic parse of the sentence in the form of an n-ary argument structure (Figure 1).

The semantic parse obtained in this manner specifies the actual action semantics for the application. One of the key derivations is the frame structure that is built upon this semantic parse. Based on the verbs occurring in the semantic parse and respective synonyms, the NLU unit constructs a frame-based semantic structure [11,12,13,14],

-------------- Tagging User Request -----------
I/PRP want/VBP to/TO invest/VB 1000/CD dollars/NNS

-------------- Chunking User Request -----------
(NX I/PRP NX) (VX want/VBP to/TO invest/VB VX) (NX 1000/CD dollars/NNS NX)

-------------- Semantic Parse of the request in the form: (Verb-Subj-Obj-Obj) -----------
("invest" "I" "1000 dollar")

Figure 1: Natural Language Parse

which is then correlated with the lexical predicates in the OMCSNET.

The frame structure comprises of hierarchical event-object structures derived from the semantic parse and chunked-text. This kind of generic type-based construction has subsequent positive implications on goal planning and iterative interaction with the OMCSNET [15].

Action Planning

The system needs to map the derived semantics from the user's utterance to the intentional goal structures and in turn to its own application level goal planning. As the investment-strategy process is iterative and complex, it seemed appropriate to model it using finite state automata, where states are characterized by the various steps needed to lay out an investment strategy and the transitions encode various choices that the user can express using natural language. Essentially, goals have slot-filler type structures and by progressing through the state automata, it is made feasible to attain the level where adequate investment advice could be extracted from the expert system. For instance, the frame structure for the "invest_goal" looks as following (Figure 2).

<<<Frame Name: invest>>>
Type : event
Subject : I
Objects:
Object 1 **:<<obj1>>**
Type : dollar
Attributes :
Attribute 1**:<<attr1>>**
Attribute_Name: individuation
Attribute_Value : 1000

Figure 2: Frame Representation for <invest>

<<<Frame Name: buy>>>
Type : event

Subject : USER

Objects:

Object 1 **:<<obj1>>**

Type : THING

Attributes :

Attribute 1:**<<attr1>>**

Attribute_Name: TEMPORAL

Attribute_Value : _time_value

Figure 3: Frame Representation for <buy>

Similarly, the invest_action requires an "invest" frame, where the slots pertaining to the investment object is filled with the money to be invested. Naturally, maintaining frame semantics of an utterance has advantages as the utterance frames can be compositionally correlated with the action frames (for example, subsumption criteria), thus providing a computationally efficient and an incremental approach to collaborative goal-oriented action planning and goal execution.

Common Sense Inference

One of the key issues that we address is that of defeasibility of computational semantics. Frame semantics is an elegant framework for characterizing fully specified semantics. However, due to the inherent ambiguity and potential for multiple senses, it becomes essential to correlate the fully specified frame semantics to the relevant senses. Also, from the action planning point of view it is necessary to articulate necessary and sufficient steps to achieve the desired goal (Figure 3). The common sense knowledge fulfills both of these requirements as it encodes multiple senses in a semantic network, where traversal along a particular path could reflect various steps needed to complete a particular goal [16].

As mentioned earlier, we use OpenMind [2] as the source for the common sense knowledge. OpenMind is a web-based collaborative project that aims towards acquiring knowledge in the form of English sentences that we use in our day-to-day activities. This knowledge is structured in a semantic network that specifies predicate-based semantics.

Therefore, we construct relevant queries pertaining to the user's goal, which is used to gather other senses of the goal as well as other goals which are required to achieve the goal. For instance, a typical OMCS query 'buy house', produces binary predicate structures:

1. (EventForGoalEvent "buy house" "apply for mortgage")
2. (EventForGoalEvent "buy house" "ask for loan")
3. (EventForGoalEvent "buy house" "avoid house with termite")
4. (EventForGoalEvent "buy house" "be careful")
5. (EventForGoalEvent "buy house" "contact real estate agent")
6. (EventForGoalEvent "buy house" "contact your local real estate agent")

Moving Into the Expert Domain

At this point we do some handholding with the user to better define the goal. We split this phase into three parts.

First, we use the concepts as queries and crawl the web to get the most relevant links that offer information about the goal. So, from the earlier example of "buy house", common sense comes back with facts like "real estate". The links that CSI returns will pertain to contacting real-estate agents and buying a house. The web provides a wealth of well-conducted research on various topics hence offering the expertise required to narrow down the goal [17, 18]. We carefully extract the best links and display it in a menu along with an in-built browser for the user to navigate.

Second, the user now navigates the web to get more information about the goal. While this is happening, our tool is "listening" to the hyperlinks (Figure 4). When the user finally closes on a price or value it is passed to the system and in the backend the current URL is captured for two reasons. One is to be able to return to the site at a later point either for debugging purposes or to redo the selection. Two is to extract other options that the tool can suggest to the user, if the current choice was not good.

Part three of our expert system is an information filter where we have an agent that goes out to the web looking for financial information particularly pertaining to making investments. The search is intelligent in that it looks at different industries and companies and extracts the factors that affect the performance of the market, like the volatility, price-earnings ratio, etc.

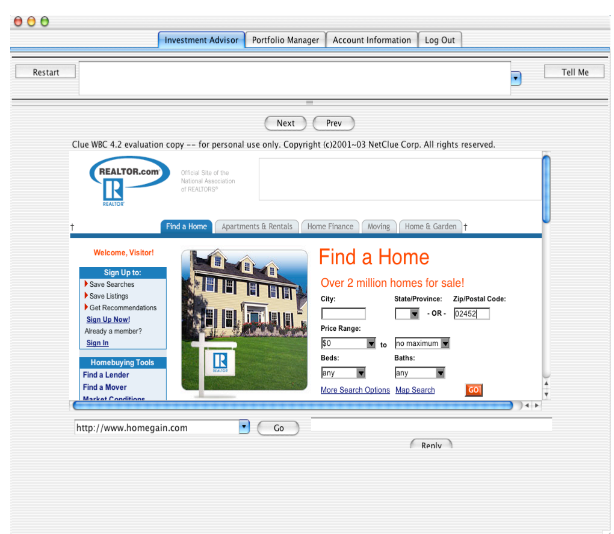

` Figure 4: An embedded WWW browser functionality

At the point this filter is triggered, the agent has at its disposal the asset-allocation determined earlier. So the input to the filter will be the expected performance of the various investments (like stocks, bonds, money market) in order to meet the user's goal and if necessary make a profit.

Common Sense Investment Strategy

The initial asset allocation is determined from the users' goal, timeline and risk tolerance. Now, we delve into each allocation and use a combination of common-sense [19] and expert knowledge to pick the top performing industries and companies the user may consider investing in and we explain the reasons behind making this selection (Figure 5).

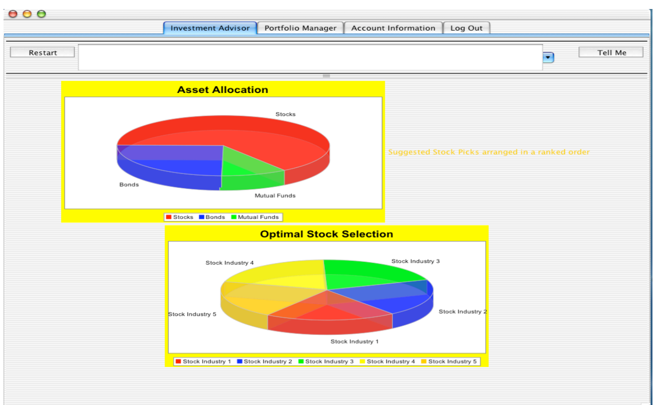

Figure 5: A Sample Asset Allocation

The typical domain specific common sense facts used are as follows:

1. ‘high risk’ -> ‘high return’
2. ‘high return’ -> ‘invest in stocks’
3. (PropertyOf "diversified stock" "good growth with high consistency over long term")
4. (PropertyOf "good stock" "larger the growth rate of dividends and earnings")
5 (CapableOf "high stock allocation" "good return for small amount of capital")

Discussions and Future Directions

CSI provides an exciting platform for investment related activities. Our effort heavily emphasizes the importance of always involving the user in the interaction howsoever tough the application domain is. The system always maintains a dialogue with the user whether it is making investment choices for him or it is trying to debug some prior investment activity. Nevertheless, it turns out that mapping from naïve knowledge to expert knowledge and vice versa is not trivial. Also, there are subtle issues related to optimality of common sense knowledge required to ascertain sufficiency for certain goal and how this can be characterized dynamically? Also, it would be interesting to see how this kind of interaction leads to social role building. As part of future activity, we aim to gather more investment-related common sense knowledge. We would also like to do better risk analysis and add functionality to the expert financial engine. For now, we have limited our allocation to just stocks, bonds and money market. This may be extended to retirement funds etc.

ACKNOWLEDGEMENTS

We would like to thank Henry Lieberman, Hugo Liu and Push Singh for numerous suggestions and encouragement.

REFERENCES

1. http://web.media.mit.edu/~hugo/omcsnet/index.html
2. http://commonsense.media.mit.edu/cgi-bin/search.cgi
3. http: //www.vanguard.com
4. http://www.fool.com
5. http://www.nasdaq.com
6. http://www.morningstar.com
7. http://www.investopedia.com
8. http://finance.yahoo.com
9. http://www.cnn.com/money
10. End-User Debugging for Electronic Commerce, (with Earl Wagner) ACM Conference on Intelligent User Interfaces, Miami Beach, January 2003.
11. Fillmore, Charles J. (1968): The case for case. In Bach and Harms (Ed.): Universals in Linguistic Theory (pp. 1-88), Holt, Rinehart, and Winston, New York.
12. Fillmore, Charles J. (1976): Frame semantics and the nature of language. In Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech, Volume 280 (pp. 20-32).
13. Fillmore, Charles J. (1982): Frame semantics. In Linguistics in the Morning Calm (pp. 111-137), Hanshin Publishing Co., Seoul, South Korea.
14. Fillmore, Charles J. (1985): Frames and the semantics of understanding. In Quaderni di Semantica, Vol 6, No. 2 (pp. 222-254).
15. Conference and Workshop Papers:Timothy Chklovski. LEARNER: A System for Acquiring Commonsense Knowledge by Analogy. In Proceedings of Second International Conference on Knowledge Capture (K-CAP 2003). October 2003.
16. The Society of Mind, Marvin Minsky.
17. Hugo Liu, Henry Lieberman, Ted Selker. (2002). GOOSE: A Goal-Oriented Search Engine with

Commonsense. In De Bra, Brusilovsky, Conejo (Eds.): Adaptive Hypermedia and Adaptive Web-Based Systems, Second International Conference, AH 2002, Malaga, Spain, May 29-31, 2002, Proceedings. Lecture Notes in Computer Science 2347 Springer 2002, ISBN 3-540-43737-1, pp. 253-263.

18. Hugo Liu and Henry Lieberman (2002). Robust photo retrieval using world semantics. In Proceedings of the LREC 2002 Workshop on Creating and Using Semantics for Information Retrieval and Filtering: State-of-the-art and Future Research, Las Palmas, Canary Islands, pp. 15-20.

19. Collecting commonsense experiences (with Barbara Barry). 2003. Proceedings of the Second International Conference on Knowledge Capture (K-CAP 2003). Florida, USA.

20. http://web.media.mit.edu/~hugo/montylingua

Hugo Liu and Push Singh (forthcoming) .OMCSNet: A Practical Commonsense Reasoning Toolkit. ***Submitted to the 2004 International Conference on Knowledge-Based Intelligent Information & Engineering Systems.***