myVU: A Next Generation Recommender System Based on Observed Consumer Behavior and Interactive Evolutionary Algorithms

Andreas Geyer-Schulz[1], Michael Hahsler[1], and Maximillian Jahn[1]

Abteilung für Informationswirtschaft,
Institut für Informationsverarbeitung und Informationswirtschaft,
Wirtschaftsuniversität Wien, A-1090 Wien, Austria

Abstract. myVU is a next generation recommender system based on observed consumer behavior and interactive evolutionary algorithms implementing customer relationship management and one-to-one marketing in the educational and scientific broker system of a virtual university. myVU provides a personalized, adaptive WWW-based user interface for all members of a virtual university and it delivers routine recommendations for frequently used scientific and educational Web-sites.

1 Introduction

In this article we describe myVU, a next generation recommender system based on observed consumer behavior and interactive evolutionary algorithms and report on first results on its usage. myVU is a prototype of an educational and scientific recommender system for providing routine recommendation and consulting services in mass universities to meet the challenge of reengineering the university posed in Tsichritzis (1999). myVU has been developed with the explicit goal of supporting life-long learning for large parts of the population and, therefore, must address the problems caused by the resulting high heterogeneity in students' background, capabilities, and previous experience. In this article, however, we emphasize the role of evolution in the design of such systems and concentrate on the evolution strategies embedded into myVU. In section 2 we critically review existing interactive evolutionary algorithms and we identify the two main obstacles for their application to web-site personalization and design. Section 3 has the aim of explaining, how evolutionary algorithms can be integrated in such a system: We treat a virtual university as an information market and base the evolutionary algorithm shaping the user interface on well-known statistics which are widely used in the retail industry, namely ABC-analysis, market-basket analysis (see Blischok (1995)) and consumer purchase histories (and – in the near future – stochastic consumer choice models, for a survey see Wagner and Taudes (1987)). These statistics are computed from appropriately instrumented Web-server transaction logs. The evolutionary operators, namely fitness-biased selection and various kinds of mutation operators, influence the design and layout of the user-interface and thus influence user behavior. Obviously, myVU can serve as a model of the business to customer interface of future e-commerce applications. myVU supports customer relationship management and one-to-one marketing as described in Kelly (1998) to a considerable extent. The evolutionary algorithm described above is common to all myVU recommender services. All myVU recommender services available today (March 2000) are presented in section 4 together with a first analysis of their usage. Finally, we describe possible future improvements of the system.

2 Interactive Evolutionary Algorithms for Design

In the opening lecture of the 1997 Genetic Programming Conference John Koza identified in his outlook on the future of evolutionary algorithms web-site design as one of the most promising (commercial) application areas. However, not much has been achieved in this area in the last three years. To be honest, we must be a little bit more precise. Interactive evolutionary and genetic algorithms have flourished in music and arts. See e.g. the web-site of Craig Reynolds (http://www.red.com/cwr/evolve.html) for examples on the WWW.

An evolutionary algorithm is characterized by a select-and-mutate approach with a population size of 1, a genetic algorithm in addition by a crossover operator and a population size larger than 1. In this paper, however, we will neglect these differences. The reader may think of an evolutionary algorithm as a border line variant of a genetic algorithm with zero probability of crossover and a population size of 1.

A survey of applications of interactive evolutionary and genetic algorithms for system design from an engineering perspective is presented in Takagi (1996a) and Takagi (1996b). However, we are not aware of any application of interactive evolutionary or genetic algorithms in the area of web-site personalization for e.g. customer relationship management and one-to-one marketing. Why?

Before we answer this question, we give a short review of the two generations of interactive evolutionary and genetic algorithms used today by artists, musicians and engineers. The main feature of the first generation of these algorithms is that the user of such an algorithm is required to **explicitely and directly** assign a fitness value to each of the objects bred by the algorithm as shown in Figure 1. The ancestor of these algorithms is Richard Dawkins' "Blind Watchmaker" which bred simple tree-based forms – so called biomorphs – with an interactive select-and-mutate approach. See Dawkins (1986). Smith's (1991) interactive genetic algorithm for breeding bug like creatures and Caldwell and Johnston's (1991) approach of assisting a witness in building a facial composite of a criminal suspect with an interactive genetic algorithm are two additional representatives of this class of algorithms. Combined with a virtual reality environment and bio-sensors as proposed by Häfner and Rößler (1995) interactive evolutionary algorithms serve as advanced industrial design environments.

The class of interactive evolutionary algorithms requires that we can phrase the problem as a search through some parameter space, that we can generate new candidate solutions in near real-time and that the utility of those candidate solutions can easily be compared by humans, but not by means of a precisely defined fitness

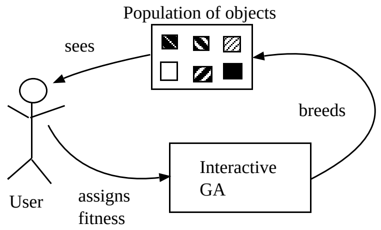

Fig. 1. Early interactive genetic algorithms.

function. Crucial for the success of this approach is that the number of candidate solutions that humans must evaluate is kept as low as possible, because baby-sitting for an interactive-evolutionary algorithm is not a task humans volunteer for.

This problem – the human fitness bottleneck – is addressed by interactive evolutionary algorithms of the second generation with the help of a meta-level genetic algorithm as shown in Figure 2.

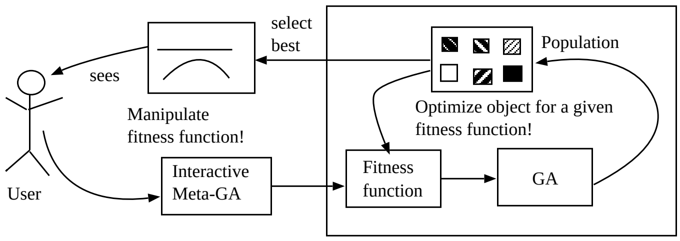

Fig. 2. Interactive (meta) genetic algorithms.

The interactive evolutionary algorithm at the meta-level breeds fitness functions which are used by a second genetic algorithm at the base-level to generate new candidate solutions. For each fitness function, only the best solution is presented to the user for evaluation. The user evolves (either implicitely or explicitely) the fitness function which captures his taste or experience or aesthetics. An example of such an interactive algorithm, where the user explicitely manipulates a set of weights which express the relative importance of design factors as structural configuration, harmony with the surrounding environment, and slenderness in the design of aesthetic bridge structures, is presented in Furuta et al. (1996). Biles et al. (1996) carry the concept one step further. In GenJam, an interactive genetic algorithm for breeding jazz-solos, the user implicitely chooses between fitness functions in the form of neural networks. By choosing a jazz-solo, he increases the fitness of the associated neural network.

However, the human fitness bottleneck can be tackled with a different kind of evolutionary algorithm which has been employed e.g. on the International Genetic Art II site by John Mount, Scott Neal Reilly and Michael Witbrock (described at http://www.cs.cmu.edu/˜jmount/g3.html) at Carnegie Mellon. In this approach, many users evaluate the objects bred by the genetic algorithm, the evaluations are collected and aggregated by the fitness function. A new generation of objects is generated as soon as a certain number of evaluations have been collected or a pre-specified period of time has passed. This approach suffers from two draw-backs: first, the user receives no immediate feedback and he can not see, how his evaluations have influenced the outcome of the algorithm and, second, if users have non-homogeneous utility functions and this is the case for heterogeneous user groups, the averaging process inherent in the aggregation of the fitness values may lead to art objects whose aesthetic is a disappointment for all users. This effect is even more pronounced with small population sizes.

The main problems for applying interactive evolutionary algorithms currently used by artists and engineers for personalized web-design are:

1. The human fitness bottleneck caused by the required explicit evaluation of designs by the users.
2. The heterogeneity of user groups which leads to non-homogeneous fitness functions.

3 The Design of myVU: Addressing the Human Fitness Bottleneck and Heterogeneous User Groups

To address the problems identified in the previous section, the design of myVU is based on the metaphor of an information market (broker) and on the principle of self-assessment of experience. This is an application of the economic principle of self-selection which has been suggested as a rationale for versioning information products by Shapiro and Varian (1999). In myVU we generate recommendations on information products (web-sites) from observed purchasing behavior for information products. The key idea is identify the information market as a system and the rest of the Internet as its environment. Whenever the user follows an external link from the information market to an information product and crosses the system boundary, this is registered as a purchase incident. Use of internal links in the broker reveals the users' preferences for broker services. Information products in the virtual university have a rich meta-data description with one or more category attributes. In addition, in myVU users have the opportunity to incrementally establish a profile of their experience with each information product category visited (self-assessment of experience).

In the following the main loop of the interactive evolutionary algorithm of myVU shown in Figure 3 is explained in more detail. Numbers in parenthesis refer to arcs in Figure 3. In myVU recommendations are usually presented as ranked and labelled lists of suggested information products (web-sites).

1. The user perceives (1) a list of labelled links to other web-sites (e.g. his personal favorites shown in Figure 4) which constitutes an element of the user interface.

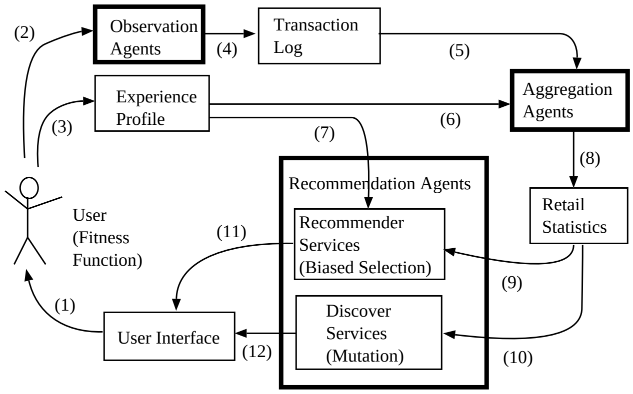

Fig. 3. myVU as interactive evolutionary algorithm.

2. By following (2) one of these links the user reveals his preference for this website. The oberservation agents record (4) the purchase incident in the transaction log. The fitness of an information product is proportional to the number of times it has been purchased. In myVU the fitness of an information product is thus not explicitely assigned by a user, it is computed from observed user behavior.
3. By incrementally revealing their experience (novice, average, advanced or expert) with information product categories they have visited in previous sessions, myVU users establish (3) their own experience profile which is used for computing group specific recommendations (6) and for selecting the appropriate group specific recommender service (7) for the user. This mechanism addresses the problem of heterogeneous user groups and has been recommended in Shapiro and Varian (1999).
4. Every night aggregation agents update the retail statistics with the purchase incidents (5) from the transaction log in accordance with the experience level of the user (6). Today, practitioners call this "web-mining". Such statistics include (depending on the degree of anonymity chosen by the user):
   (a) If we can extract anonymous user sessions from the transaction log, market baskets can be analyzed. We compute for each information product y_i the frequencies that, if information product y_i is in a market basket, the other information products $y_1, \ldots, y_{i-1}, y_{i+1}, \ldots y_n$ are also in the basket. This frequency is proportional to the conditional probability $P(y_j|y_i)$ to buy information product y_j, if information product y_i has been purchased in the same session.
   (b) If we can extract pseudonymous user sessions from the transaction log, we can establish the purchase histories of users and combine these with their experience profiles. (Note, that pseudonymous means, that we do not know the true identity of a user, we only know that it is the same user.) We compute for all information products in a category the frequencies with which they are bought by a novice, an average user, an advanced user, and an

Fig. 4. Favorite Entries: A myVU-Service.

expert. Again, this is proportional to the conditional probability $P(x_c|e_c)$ that an information product x_c from category c is purchased, if a user has experience level e_c for category c.

(c) For each user, the number of purchases of each information product in the relevant part of his purchase history. This statistic is used for ranking the entries of the *Favorite Entries* service shown in figure 4.

(d) For each user, the number of purchases aggregated for information product categories.

(e) All products bought at least once by some user.

(f) From all transactions, an ABC-analysis of information products and categories (for labelling the highest rated group with HOT).

5. As a user requires an information service of myVU, the recommendation agent responsible for this service generates the user interface element which results from this service. We distinguish between “recommender services” which im-

plement fitness-proportional selection operators of the evolutionary algorithm and "discover services" which implement mutation operators of the evolutionary algorithm. *Favorite Entries* shown in Figure 4 is an example of such a recommender service. Note, that a user interface element may consist of a bundle of such services.

According to Holland (1975) fitness-biased selection operators serve the exploitation of information. In myVU fitness-biased selection operators produce lists of information products ranked by their fitness. Ranking information products has a strong influence on the choice behavior of users. This fact has been established by several studies of user choice behavior in online environments. For a survey, see Introna and Nissenbaum (2000). In myVU, fitness-biased selection operators make repeat-buying behavior easier and thus increase user convenience. Convenience for repeat-buying has been identified as one of the key success factors for e-commerce sites by Bellmann et al. (1999).

Mutation operators have the role of supporting the exploration of information. In myVU mutation operators either have the form of randomly generated link or category lists or of a randomly drawn banner leading to an information product in the virtual university. In the current implementation of myVU the mutation operators draw from the list of all information products ever purchased by some user and from the list of all information product categories ever used by some user. However, in future releases of myVU additional mutation operators drawing from other neighborhoods (e.g. all information products and categories in the virtual university, the conditional probability distribution of cross-selling, ...) will be investigated. Mutation addresses an incentive problem of recommender systems discussed in Resnick and Varian (1997), namely, that users receiving recommendations diminish their search effort for information products and increasingly rely on a very narrow set of information products.

Consider, for example the *Favorite Entries* recommendation service illustrated in Figure 4. The link list is ranked according the the user's personal purchase frequencies computed from his purchase history. Klicking on the "link" (the recommender service) with the label `::others also use` in the line below the link "Genetische Lernverfahren" (Genetic Machine Learning) leads to a list of information products $y_1, \ldots, y_n$ ranked according to the conditional probability $P(y_j|$"Genetische Lernverfahren"$)$. The "link" (the recommender service) with the label `::experts also use` indicates that the user is an expert in the field of Genetic Machine Learning and it leads to a list of information products of the same category ranked according to the probability of being purchased by other experts for this category.

Recommendations based on the experience level of users for a category of information products are only available for users who have revealed their experience for this category. This is a tit-for-tat strategy which adresses the free-riding problem inherent in recommender systems and it offers an incentive to the user to reveal his self-assessment of his experience for a category of information products. In the future, we expect to exploit this information for learning progress monitoring and for team-building.

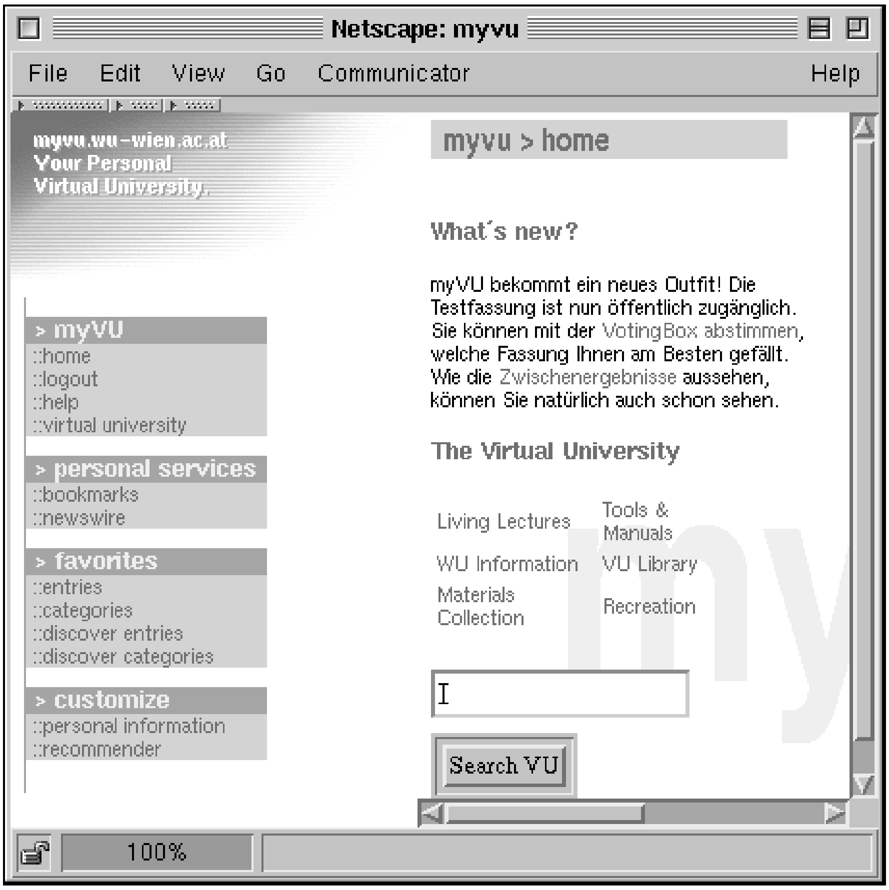

Fig. 5. myVU-Services.

4 The myVU Recommender Services

Figure 5 shows the main page of the second version of myVU. The following recommender services are available at the moment in myVU:

- *Favorite Entries* (`::entries` in the sidebar of Figure 5) have been shown in Figure 4 and discussed in the previous section.
- *Favorite Categories* (`::categories` in the sidebar of Figure 5) are a list of a user's most frequently used product categories. `::other also use` in the context of a category is based on the conditional probability distribution that a user who buys a product from one category also buys a product from another category.
- *Discover Entries* (`::discover entries` in the sidebar of Figure 5) is a list of randomly selected products from the virtual university for the purpose of stimulating the curiosity of myVU users.

Service	Usage (absolute)	Usage (relative)
Bookmarks	628	38.01%
Newswire	77	4.66%
Favorite Entries	350	21.19%
Favorite Categories	228	13.8%
Discover Entries	74	4.48%
Discover Categories	79	4.78%
Recommender Profile	146	8.84%
Personal Information	70	4.24%
Total	1652	100%

Table 1. Usage of myVU Services (January 26th, 2000 – March 13th, 2000)

- *Discover Categories* (`::discover categories` in the sidebar of Figure 5) is similar to *Discover Entries*, only at the level of information product categories. The rationale for this service is to encourage the exploration of information product categories which are new for a myVU user.

They are grouped under the heading `> favorites` in the sidebar of the myVU main page shown in Figure 5. Note, that recommender services are embedded as labels into all myVU services including the myVU global bookmark service (`::bookmarks`) which myVU users can access from wherever they are on the Internet.

In addition, the service `::recommender` under the heading `> customize` allows the user to change his level of experience for categories he has visited in a previous session in the virtual university.

In Table 1 we have listed the actual usage of myVU services in the first six weeks of operation (from January 26th, 2000 to March 13th, 2000). We see that all recommender services are actually used by myVU users. *Favorite Entries* is the most popular recommender services. It accounts for 20 percent of myVU service usage. Both mutation-based *Discover* services account for 4 percent of myVU usage each. Although this is only a very preliminary result, this seems to indicate that the mutation-based *Discover* services stimulate user curiosity and lead to an increase in the exploration efforts of users. Informally (e-mails and chat), student response to myVU has been favorable. At the end of March 2000 more than 200 users were registered myVU users.

5 Future Research

As usual, a lot remains to be done. From the perspective of evolutionary computation which we have emphasized in this article, the following problems merit further investigation:

- What is a useful concept for a crossover operator for web-site design?
- A long term and more detailed study of the use of recommender services in myVU is required.

- Several (obvious) refinements of the current services are not (yet) implemented: cut-off values for truncating probability distributions, discounting strategies for older purchases, more flavors of mutation operators, ...
- Appropriate bundling and naming of services.
- Improved statistics of the dynamic development of the recommender system for an improved management of such systems.
- A study of user acceptance. Are the changes in the myVU user interface unobtrusive enough, so that the user accepts an adaptive user interface?
- What is the potential of combining recommendations based on behavior with content-related recommendations?

6 Links

We invite the reader to visit myVU (http://myvu.wu-wien.ac.at) at the Virtual University (http://vu.wu-wien.ac.at) of the Vienna University of Economics and Business Administration and to try the personalized recommender services described in this article.

7 Acknowledgement

We gratefully acknowledge the financial support of the Jubiläumsfonds of the Austrian National Bank under Grant No. 7925 without which myVU would never have been reality.

References

BELLMANN, S., LOHSE, G.L., and JOHNSON, E.J. (1999): Predictors of Online Buying Behavior. *Communications of the ACM, 42(12),* 32–38.

BILES, J.A., ANDERSON, P.G., and LOGGI, L.W. (1996): Neural Network Fitness Functions for a Musical IGA. In: *Procs. of the Int. ICSC Symp. on Intelligent Industrial Automation (IIA'96) and Soft Computing (SOCO'96),* ICSC Academic Press, Reading, B39–B44.

BLISCHOK, T.J. (1995): Every Transaction Tells a Story. *Chain Store Age Executive, 71(3),* 50–62.

CALDWELL, C. and JOHNSTON, V.S. (1991): Tracking a Criminal Suspect through "Face-Space" with a Genetic Algorithm. In: R.K. Belew and L.B. Booker (Eds.): *Procs. of the 4th Int. Conf. on Genetic Algorithms (ICGA'91),* Morgan Kaufmann, San Mateo, 416–421.

DAWKINS, R. (1986): *The Blind Watchmaker.* Longman, Harlow.

FURUTA, H., MAEDA, K., and WATANABE, E. (1995): Application of Genetic Algorithm to Aesthetic Design of Bridge Structures. *Microcomputers in Civil Engineering, 10,* 415–421.

HÄFNER, U. and RÖSSLER, A. (1995): Interactive Optimization with Evolutionary Strategies as a Tool for Virtual Reality. In: M. Freter (Ed.): *Procs. Virtual Reality World'95,* IDG, München, 369–378.

HOLLAND, J.H. (1975): *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor.

INTRONA, L. and NISSENBAUM, H. (2000): Defining the Web: The Politics of Search Engines. *IEEE Computer, 33(1),* 54–62.

KELLY, K. (1998): *New Rules for the New Economy – 10 Radical Strategies for a Connected World*. Viking Penguin, New York.

RESNICK, P. and VARIAN, H.R. (1997): Recommender Systems. *Communications of the ACM, 40(3),* 56–58.

SHAPIRO, S. and VARIAN, H.R. (1999): *Information Rules: A Strategic Guide to the Network Economy.* Harvard Business School Press, Boston.

SMITH, J. (1991): Designing Biomorphs with an Interactive Genetic Algorithm. In: R.K. Belew and L.B. Booker (Eds.): *Procs. of the 4th Int. Conf. on Genetic Algorithms (ICGA'91),* Morgan Kaufmann, San Mateo, 535–538.

TAKAGI, H. (1996a): System Optimization Without Numerical Target. In: M.H. Smith, M.A. Lee, J. Keller, and J. Yen (Eds.): *Procs. of the 1996 NAFIPS Conference*, IEEE Service Center, Piscataway, 351–354.

TAKAGI, H. (1996b): Interactive GA for System Optimization: Problems and Solution. In: H.J. Zimmermann (Ed.): *Procs. 4th EUFIT Conference (EUFIT'96),* Verlag Mainz, Aachen, 1440 – 1444.

TSICHRITZIS, D. (1999): Reengineering the University. *Communications of the ACM, 42(6),* 93-100.

WAGNER, U. and TAUDES, A. (1987): Stochastic Models of Consumer Behaviour. *European Journal of Operations Research, 29(1),* 1–23.