

EvoShelf: A System for Managing and Exploring Evolutionary Data

Timothy Davison¹, Sebastian von Mammen¹, and Christian Jacob^{1,2}

¹ Dept. of Computer Science, Faculty of Science

² Dept. of Biochemistry & Molecular Biology, Faculty of Medicine
University of Calgary, Canada

{tbdaviso,s.vonmammen,cjacob}@ucalgary.ca

Abstract. Systems that utilize evolutionary computation produce large amounts of data. Quite often, this data has a convenient visual representation. However, managing and visualizing evolutionary data can be a difficult and onerous task. By employing techniques used in photo management software, we have produced a system that helps to organize and visualize evolutionary data while retaining a complete record of a simulation. By means of a simple plugin architecture this system can be extended to import data produced by arbitrary evolutionary systems. We present the system's architecture, its features, and we provide a comprehensive example, highlighting its advantages in applied research.

1 Introduction

Evolutionary systems produce large amounts of data. Beyond the obvious data (such as the genotype and phenotype of an individual), there is a considerable amount of meta-data produced as well. Such data includes the hereditary data, fitness values, and other attributes of the evolutionary computation approach being employed.

It is common to manage experimental data by means of a file-system browser, such as the Finder in Mac OS X, and Windows Explorer in Microsoft Windows. Searching or organizing individuals according to various criteria is a laborious task in such systems. Consider a system that organizes the individuals produced by an experiment into sub-directories by generation, giving each individual its own file containing its genotype, and phenotype, along with meta-data such as fitness, or genealogy. Filtering these individuals by fitness value would be a difficult task with either file-system browser.

An evolutionary system may employ an interface of its own for browsing the data that it produces. In this case, the visualization procedures and the management of the genotype/phenotype data are typically implemented specifically for the one evolutionary system. However, the universality of evolutionary algorithmic approaches renders generic visualization and data management techniques valuable across various application domains.

In a way, the situation is very similar to managing individual (digitized) image and music collections. Such libraries can easily consist of thousands of items. A

number of applications have made the organization and management of such data much easier for the end user [1,2,4,9]. We propose a system that can deal with evolutionary data with the same ease of use and flexibility as provided by these mainstream media management applications.

EvoShelf is an extensible system that allows for the importing and exploration of evolutionary data from evolutionary systems. It solves the challenge of organizing imported metadata by providing a navigable, and searchable image-based browser that uses interface design elements from Apple's photo and music management software iPhoto [1], and iTunes [2]. Furthermore, it provides a plugin framework for building additional import modules and visualizations. Both navigation and visualization are optimized for real-time interaction.

In Section 2 we explore the topic of visualizing data in evolutionary systems. Section 3 presents the design of the *EvoShelf* system and its graphical user interface. It also touches upon some of the example visualization techniques included in *EvoShelf*, as well as details about its plugin architecture. In Section 4 we use *EvoShelf* in coordination with an existing evolutionary system for semi-interactive evolutionary computing, and for analyzing the results produced by that system. We will conclude in Section 5 with a summary of our work, along with possible directions in which to take it in the future.

2 Related Work

We briefly outline the data management and user-interface approach of various software systems that inspired the *EvoShelf* visualization and management system. Secondly, we outline various techniques that have been developed for visually supporting computational evolutionary experiments.

2.1 Digital Media Libraries

The framework presented in this article was mainly inspired by iPhoto, Apple's mainstream photo management application [1]. It is capable of organizing and browsing thousands of photos. Despite the large amounts of information that it is capable of presenting to the user, it maintains a very simple and intuitive interface. It consists of two primary views, an organizer view, and an image browsing view (Figure 1(a)). Multiple images, up to and including an entire library of photos, are displayed in the browsing view. The organizer view is used to filter this view into subsets of photos, such as those represented by a photo album containing the user's favorite photos. As photos are imported into the system they are grouped into events. Pictures taken during a certain period of time might have all been taken during a vacation and the respective group of photos could be labeled after the location of the recreational stay.

iTunes is another application from Apple Inc. that manages a large amount of data in a similar fashion to iPhoto. Unlike iPhoto, whose interface is focused on visualizing and managing photos, iTunes is targeted towards playing music and organizing large digital music collections. Visual cover art often decorates

individual music files, but the iTunes library is mainly organized by sorting and searching through textual meta-data such as artist name, music category, or album name (Figure 1(b)). Together, iPhoto and iTunes suggest an interface that combines visualization and meta-data management techniques that could be very powerful for organizing evolutionary data.



Fig. 1. The user interfaces of the media management applications (a) iPhoto and (b) iTunes

2.2 Evolutionary Visualization Techniques

Various data visualization techniques have been presented in the context of evolutionary computing. On the one hand, individuals can be compared at a glance based on their multi-dimensional genotypes, independent of the respective interpretation or phenotype. On the other hand, methods of visualization have been developed that capture characteristics of whole populations, allowing one to visually track the evolutionary process.

Pohlheim, for instance, presented a toolkit of convergence diagrams, 3D line plots, and 2D image plots, to visualize the evolution of fitness values and other individual attributes [10]. Hart and Ross introduced a tree-based visualization to trace the ancestry of the best individual produced by an evolutionary run [5]. Daida et al. unfold genetic ancestry onto concentric circles on a 2D plane to create a compact and highly scalable visualization of hereditary processes [3]. Wu et al. represent genotypes as sequences of color coded stripes whose colours correspond to different genes [12]. Keim et al. designed a system to visualize search queries on a (relational) database [7]. Data items that match the query most closely are arranged in the center of a spiral arrangement. This visualization technique can be used to relate individuals in an evolutionary system in arbitrary ways, e.g. by comparing fitnesses or individual attributes. In [8], Khemka and Jacob have closely investigated the possibilities to visualize population-based

optimization processes at various levels of scale—from the individual to sets of experiments. They provide an easily adaptable user interface with various interactive manipulators to explore optimization processes across these scales.

3 The *EvoShelf* System

The interface of *EvoShelf* is divided into three window panes (Figure 2). The organization view on the left-hand side is used for selecting and grouping imported experimentation data (Section 3.1). The user’s selection is shown in the browser view in the center pane. An inspector view (right-hand side) shows further details about an individual or an experiment. In addition to importing and inspecting functions, the toolbar at the top of the window gives access to built-in visualization methods which are explained in Section 3.2. Typically, a user of *EvoShelf* writes a plugin to import and visualize data for his respective evolutionary system, if it does not already exist. We provide details about plugins in Section 3.3.

EvoShelf makes use of lazy fetching of data. That is, images and attributes of an individual are only loaded when they are needed. When individuals go off screen, their data is unloaded. In this way, we have manipulated data sets with over 40,000 individuals. Conceivably, *EvoShelf* can work with even larger datasets. To further increase the scalability of *EvoShelf*, high resolution images of individuals are loaded on demand—if no zoom is required, a low resolution image is displayed instead.

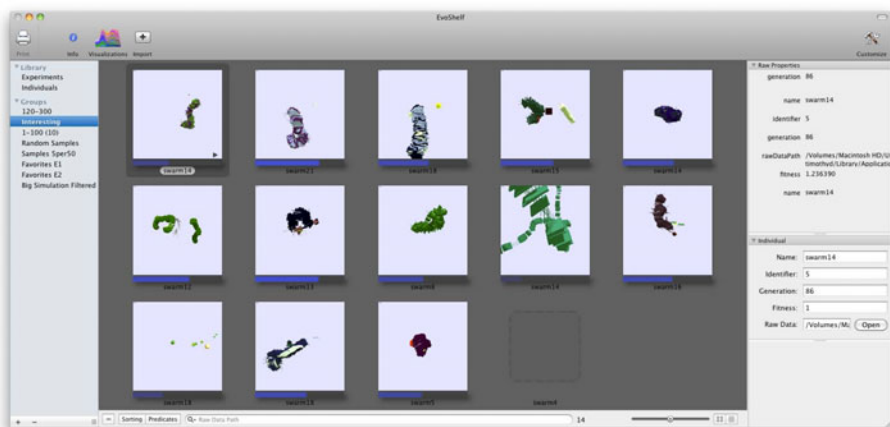


Fig. 2. The graphical user interface of *EvoShelf*

3.1 Individuals, Experiments, and Groups

The organizational view to the left of Figure 2 is divided into a library section and a groups section. In the library section, the user can select either *Individuals*

or *Experiments*. In particular, the *Individuals* selection displays the images of all the individuals in the library, whereas *Experiments* shows representative thumbnails of all the imported experiments. The user can browse through the set of individuals of any experiment by hovering with the mouse over its thumbnail. The individuals of an experiment are revealed when the user double clicks on the experiment.

The data in the browser view can be sorted or filtered by the experiments' and individuals' attributes. Once the user has formed a suitable selection he can save his selection in a group, which would be equivalent to photo albums or playlists in [1,2] respectively. In Figure 2, a group labelled *Interesting* is selected, which hosts individuals from multiple experiments that the authors found of interest. Groups can be organized hierarchically. That is, one can form groups containing groups. When such a group is selected a union is formed from all of the individuals contained within the subgroups.

The controls at the bottom of the interface allow the user to remove individuals from a group or from the library, to sort individuals, to search for individuals according to arbitrary attributes (such as fitness value or generation), to scale the size of the images displayed, and to change the display mode. One display mode shows individuals as a collection of images, another one lists them in tabular format. The latter view is convenient for sorting and searching through individuals based upon numeric or textual attributes.

The specimen in the upper left corner is selected in the browser view in Figure 2. The image representing the individual was generated by the evolutionary system used as a test run for *EvoShelf* (see Section 4). A play button (a right pointing arrow) is projected on top of the specimen's description. It allows one to re-run the simulation that produced and/or evaluated the selected individual. If the plugin does not support re-runs of individual simulations, the play button is not shown.

Below the images in the browser view in Figure 2, blue bars represent the individuals' fitnesses. The bars are scaled to the minimum and maximum fitness of all the individuals currently displayed in the browser view. The higher the relative fitness, then the brighter and longer the individual bar is. No image is provided for *Swarm35* indicating that the genotype data was successfully imported but no image was found—in the given case, the simulation was terminated before a screenshot would have been taken.

The inspector view on the right hand side displays several default properties about the imported data, such as the file name of an individual or experiment. A custom interface for the inspector can be defined via the plugin architecture (Section 3.3).

3.2 Built-in Visualization Techniques

EvoShelf employs two basic built-in visualization techniques: star plots of name-value pairs [8] and FitnessRiver, a derivative of the ThemeRiverTM method, which integrates local numeric values with global trends [6].

A star plot in *EvoShelf* visualizes a set of name-value pairs as a series of radially arranged line segments (Figure 3(a)). The length of a line segment is representative of an attribute's value and it is normalized to the attribute's maximum value in respect to the selected individuals. An attribute's line segment will consistently appear at the same location in a star plot to render individuals comparable.

The ThemeRiverTM visualization method produces a stream diagram that is read from left to right. Currents in the stream represent individual themes that occur, grow and decay over time. Currents are visually differentiated from each other by way of colour, and those colours may be reused for non-adjacent currents. Instead of separating equivalent attributes into individual currents of a stream diagram, our FitnessRiver visualization method stacks the fitness values of individuals on top of each other. The width of a current is proportional to the fitness of the corresponding individual. Different colours are used to distinguish between individuals. Discontinuing currents indicate the removal of an individual from the evolutionary process. In the FitnessRiver visualization the x-axis represents the sequence of generations. A flat baseline is used so that the user has a greater sense of the progression of the fitness evolution (Figure 3(b)).

In Figure 3(b) we can see a large jump in the overall fitness at about the middle generation. When we look closely, we see that there are a few very successful individuals in the previous generation. We can see how these individuals likely contributed to the next generation. Furthermore, the majority of individuals in the new generation have noticeably more fitness than those in the previous generation.




Fig. 3. (a) Individuals are comparable based on their star plots. (b) The FitnessRiver visualization shows the evolution of local and the global fitness. It is an adaptation of ThemeRiverTM [6].

3.3 Plugins

A user can define additional import modules, visualization modules, data models, and finally custom inspector views for custom data models¹. A few basic

¹ *EvoShelf* plugins are written in Objective-C and should follow Apple's Cocoa API <http://developer.apple.com/cocoa/>.

classes are provided for these modules and models that serve as plugin templates. The importing process, including control over import dialogue windows, can be adapted and alternative visualization modules can be subclassed from the *EvoShelf* visualization view controller class.

The default data model (Figure 4) is well suited for generational systems such as evolutionary algorithms (EA) but also supports other heuristic computation concepts, such as particle swarm optimization (PSO). For instance, each step in a PSO simulation could correspond to a generation in an EA. *EvoShelf* can be adapted for other evolutionary systems by extending the default data model to add new attributes, or relationships. For instance, the set of attributes of the classes *EVExperiment* and *EVIndividual* can be adapted to match a given evolutionary system. The new attributes automatically determine the searching and sorting options in *EvoShelf*, as well as the information provided by the inspector view. In case a more elaborate inspector view is desired, an interface constructed in Apple’s WYSIWIG Interface Builder application can be loaded.




Fig. 4. The default data model for importing and managing *EvoShelf* data

4 Example Scenario

In this section, we explore the use of *EvoShelf* with a preexisting evolutionary system. In the evolutionary system of choice, Swarm Grammars (SGs) are bred by means of a Genetic Programming algorithm to produce architectural idea models [11]. SGs are a swarm-based developmental model in which production and interaction rules guide the movements, constructions and the reproduction of agents in 3D space.

In a subdirectory for each generation, the genotypes are stored as text files and snapshots of the corresponding phenotypes as images. Fitness evaluations for the individuals are stored in an additional file. When importing all the individuals, including their image representations and their meta-data into *EvoShelf*, the original directory structure is automatically copied into *EvoShelf*’s database.

Figure 5(a) shows a set of interesting SG specimens. We want to emphasize that due to the partially very low fitness values of the 2nd, 4th and 20th swarms from the top, we would have very likely not have inspected these phenotypes without relying on *EvoShelf*’s visual browsing functionality. Based on these undervalued, interesting phenotypes, we were able to improve the fitness function that drives the SG evolution by repositioning the geometrical focus of the fitness evaluation in respect to the SGs’ constructions to better suit the favored ones. We also used *EvoShelf* for a semi-interactive evolutionary process by repeatedly selecting and exporting interesting individuals, modifying the fitness function and parameters to the GA, breeding their offspring for a fixed number of generations and importing the outcome (Figure 5).




Fig. 5. (a) 20 interesting individuals are selected from an experiment and served as the initial generation for a (b) follow-up experiment

We discovered that the SG GP evolution usually converged prematurely after at most several hundred iterations. Figure 6 shows the FitnessRiver plot over 300 generations. Overall 20,000 individuals were computed and imported into *EvoShelf*. We noticed that the overall fitness of our individuals had stagnated by the 100th generation (there is a very slight improvement in fitness past this point). Figure 7 confirmed our assumption of over-fitting: Up to the fitness stagnation at around generation 100, we randomly chose and plotted one of the ten best individuals every ten generations. For the period afterwards, we plotted one of the ten best individuals at random every 20 generations. And indeed, the phenotype images in combination with the star plots reveal a one-sided development, most easily recognizable by the inverted T-shaped star plots. Upon closer investigation, this similarity corresponds to the deployed amounts of two out of three construction elements provided to the SG agents (rods and layers), and the amount of construction elements that were placed outside of the intended target area. As the latter construction elements reduce the fitness of an individual, their increase might explain the fitness fluctuation as observed in Figure 6.




Fig. 6. The FitnessRiver plot shows stagnating and fluctuating fitness development after about 100 generations. The vertical lines denotes each 50th generation.



Fig. 7. First, for every ten generations, then (2nd half) for every 20 generations, a star plot and phenotype of a randomly selected individual is shown

5 Summary and Future Work

We presented *EvoShelf*, an easy-to-use application in the same vein as mainstream media-browsers for managing the experimental data produced by arbitrary evolutionary systems. Fast browsing of supplementary images associated with each specimen or of generic visualizations of those specimens enables the user to retrace and interactively explore vast amounts of data. Storing, retrieving, and ordering experimental data is facilitated by a simple yet powerful search function that considers a specimen's attributes and meta-data (generation, fitness, etc). Hierarchical grouping structures further facilitates the management of large data sets. In addition to the built-in management and visualization methods, *EvoShelf*-can be extended with plugins that implement new import, visualization or inspection functionalities. According programming templates are provided that can be easily adjusted or majorly extended, depending on the user's demands.

We applied *EvoShelf* to an evolutionary application that breeds Swarm Grammars to generate architectural idea models [11]. Due to the convenient and fast browsing functionality of *EvoShelf*, we have been able to identify specimens that

received low fitness values despite their appeal. As a consequence, *EvoShelf* helped us to adjust the fitness function of the SG GP system to better suit our expectations. By means of the visualization techniques that come with *EvoShelf*, FitnessRiver and star plots, we have been able to track and investigate an over-fitting process in our evolutionary runs. Finally, by using the selection and storing capabilities of *EvoShelf*, we have also been able to introduce interactivity into an otherwise autonomous evolutionary process.

In the future, we would like to add more visualizations, and to improve the current ones. For instance, it should be possible to automatically overlay different individual-based visualizations as we have done in Figure 7. The star plot visualization should be extended to improve its readability—possibly by using different coloring schemes, or line strengths. Overall, we found it would be useful to automatically associate representative specimens with global trends, as attempted by the combination of Figures 6 and 7 or by revealing details as one mouses over a visualization. The FitnessRiver visualization could also use new visual queues to track the application of genetic operators and the course of inheritance. We would also like to explore the possibility of using *EvoShelf* with running systems, for controlling systems that use interactive evolution as found in [8], and to directly execute and manage experiments. Finally, we would like to release the system as open source software.

References

1. Apple Inc. Apple - iPhoto (April 2010), <http://www.apple.com/ilife/iphoto/>
2. Apple Inc. Apple - iTunes (April 2010), <http://www.apple.com/itunes/>
3. Daida, J., Hilss, A., Ward, D., Long, S.: Visualizing tree structures in genetic programming. *Genetic Programming and Evolvable Machines* (January 2005)
4. Google. Picasa photo editing (April 2010), <http://picasa.google.com/>
5. Hart, E., Ross, P.: Gavel—a new tool for genetic algorithm visualization. *Evolutionary Computation* (January 2001)
6. Havre, S., Hetzler, B., Nowell, L.: Themeriver tm: In search of trends, patterns, and relationships. *IEEE Transactions on Visualization and Computer Graphics* (January 2002)
7. Keim, D., Kriegel, H.: Visdb: Database exploration using multidimensional visualization. *IEEE Computer Graphics and Applications* (January 1994)
8. Khemka, N., Jacob, C.: Visplore: a toolkit to explore particle swarms by visual inspection. In: *GECCO 2009: Proceedings of the 11th Annual Conference on Genetic and Evolutionary computation* (2009)
9. Nullsoft. Winamp media player (April 2010), <http://www.winamp.com/>
10. Pohlheim, H.: Visualization of evolutionary algorithms - set of standard techniques and multidimensional visualization. In: *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO* (1999)
11. von Mammen, S., Jacob, C.: Evolutionary swarm design of architectural idea models. In: *GECCO 2008: Proceedings of the 10th annual Conference on Genetic and Evolutionary Computation* (July 2008)
12. Wu, A., Jong, K., Burke, D., Grefenstette, J., Ramsey, C.: Visual analysis of evolutionary algorithms. In: *Proceedings of the 1999 Congress on Evolutionary Computation, CEC 1999* (1999)