


available at [www.sciencedirect.com](http://www.sciencedirect.com)journal homepage: [www.elsevier.com/locate/ecolmodel](http://www.elsevier.com/locate/ecolmodel)

# The use of the GARP genetic algorithm and Internet grid computing in the Lifemapper world atlas of species biodiversity

David R.B. Stockwell<sup>a,\*</sup>, James H. Beach<sup>b</sup>, Aimee Stewart<sup>b</sup>, Gregory Vorontsov<sup>b</sup>, David Vieglais<sup>b</sup>, Ricardo Scachetti Pereira<sup>c</sup>

<sup>a</sup> San Diego Supercomputer Center, 9500 Gilman Drive, La Jolla, CA 92037, United States

<sup>b</sup> Biodiversity Research Center, 1345 Jayhawk Boulevard, University of Kansas, Lawrence, Kansas 66045, United States

<sup>c</sup> Centro de Referencia em Informacao Ambiental, Av. Dr. Romeu Tortima, 388, Campinas, Sao Paulo, 13084520, Brazil

## ARTICLE INFO

### Article history:

Published on line 20 December 2005

### Keywords:

GARP

Lifemapper

Species distribution

Genetic algorithm

Biodiversity

## ABSTRACT

Lifemapper (<http://www.lifemapper.org>) is a predictive electronic atlas of the Earth's biological biodiversity. Using a screensaver version of the GARP genetic algorithm for modelling species distributions, Lifemapper harnesses vast computing resources through volunteers PCs similar to SETI@home, to develop models of the distribution of the world's fauna and flora. The Lifemapper project's primary goal is to provide an up-to-date and comprehensive database of species maps and predictive models (i.e. a fauna and flora of the world) using available data on species' locations. The models are developed using specimen data from distributed museum collections and an archive of geospatial environmental correlates. A central server maintains a dynamic archive of species maps and models for research, outreach to the general community and feedback to museum data providers. This paper is a case study in the role, use and justification of a genetic algorithm in development of large-scale environmental informatics infrastructure.

© 2005 Elsevier B.V. All rights reserved.

## 1. Introduction

Humans have explored the life of the planet for the past 250 years. That knowledge is documented by millions of original specimens of plants and animals in the world's natural history museums and herbaria. However, we have yet to achieve the goal of identifying and mapping the distribution of the species present on earth. The vision of the Lifemapper system was to use the Internet to retrieve records of species locations from museum collections, to compute the ecological profile of each species and predict where each species could potentially live. This would then be the basis of a database of species distribution maps. The approach of developing models of species'

habitat provides not only maps of where Earth's species of plants and animals live, but also predictions of where Earth's species of plants and animals could potentially live under different scenarios such as climate change, and where and how introduced species could spread across different regions of the world.

The primary data for species distribution analysis are species location records. These localities for millions of plants and animal specimens are recorded in the world's natural history museums. There is a growing recognition of the usefulness of these records for spatial modelling of biodiversity (Peterson and Stockwell, 2002). These data have also become easily accessible via the integration of the world's museum

\* Corresponding author. Tel.: +1 858 8220942; fax: +1 858 8223631.

E-mail address: [davids@sdsc.edu](mailto:davids@sdsc.edu) (D.R.B. Stockwell).

0304-3800/\$ – see front matter © 2005 Elsevier B.V. All rights reserved.

doi:10.1016/j.ecolmodel.2005.11.016

collection into a distributed database called the Species Analyst (Vieglais et al., 1998).

Developing maps of species distributions using multivariate models of species occurrence points with environmental variables is now a widespread practice in biodiversity science (e.g. Scott et al., 2002). The primary challenge of biodiversity mapping is to develop accurate maps based on statistical relationships rather than actual observations, particularly using ad hoc museum collections data rather than controlled survey data. The Lifemapper team has chosen to use an evolutionary algorithm called the genetic algorithm for rule-set production (<http://www.biodi.sdsc.edu>), or GARP (Stockwell and Noble, 1992; Stockwell and Peters, 1999; Stockwell, 1999). In comparison with other algorithms, the GARP algorithm has shown to be the best available method for reliable species predictions using small sets of ad hoc data typically returned from museum databases (Peterson and Stockwell, 2002; Stockwell and Peterson, 2002a,b). The robustness of the algorithm has contributed to its use in projects requiring the development of large numbers of species distribution models, such as change in ecological communities due to climate change (Peterson et al., 2002). One of the purposes of this paper is to describe how the qualities of the GARP genetic algorithm contribute to the success of the project.

The second main challenge of developing a fauna and flora for the world is the computational scale of the project. This problem is computationally intensive in many aspects:

- The number of species in the world that could be mapped, e.g. for birds alone already 9600 species (Sibley and Monroe, 1990).
- The fine scale of maps that could be produced, e.g. a map of the world at a resolution of 1 km at the equator has approximately 1,000,000,000 cells.
- The number of replicates that need to be produced for each species for estimating statistical variance.
- The need for computing distributions under alternative scenarios such as climate change and geographic invasion.
- The need to recompute the maps when new data becomes available.

The need to maintain an up-to-date resource by recomputing maps when new data becomes available ensures a perpetual demand for computing resources. The combination of millions of idle computers around the world connected to the Internet forms the infrastructure for Internet grid computing which seeks to exploit otherwise idle workstations, PCs and bandwidth to create powerful distributed computing systems. Internet computing was popularized by the SETI@home project that enlisted personal computers to analyze data for indications of extraterrestrial intelligence. As SETI@home is now running on half a million PCs and delivering 1000 CPU years per day, it is currently the fastest (admittedly special purpose) computer in the world.

The goal of the Lifemapper project is to become a major component in the biodiversity informatics infrastructure. The archive of maps and models for the world's species developed through Lifemapper will be an invaluable resource for researchers in terrestrial, marine and freshwater environments. This paper describes the Lifemapper resource with

particular reference to the genetic modelling algorithm that serves as the core computational component of the screen-saver modelling program.

## 2. Methods

The general approach to producing maps of species' distributions is to develop a multivariate statistical model of known species' occurrence records and environmental variables. The species occurrence data are gathered from a number of biological collections housed at several museums and herbaria worldwide. Those institutions have their specimen databases linked and integrated through the Species Analyst (Vieglais et al., 1998) project. The environmental information is composed of a set of global geographic coverages, called environmental layers. Each layer represents one particular environmental parameter, such as temperature, rainfall, land use and elevation, among others. The layers are continuous grids, where each cell contains the value of an environmental parameter at a location.

A model for estimating the probability of occurrence of a species is developed from the species occurrences and the environmental variables. Using the values of the environmental variables at each grid cell to predict the probability of the occurrence of the species over the entire grid then produces a map of the distribution of the species. An example is given in Fig. 1.

The Lifemapper system is composed of components for harvesting species occurrence data from museum databases, developing models via the screen-saver program, storing the results in an archive, and enabling access to information and membership management functions. The schematic diagram for Lifemapper is shown in Fig. 2. This is described in three stages, pre-processing, processing and post-processing.

### 2.1. Pre-processing

The Lifemapper program retrieves georeferenced data from a network of biodiversity databases that are accessible over the Internet. At this time, almost all sources utilize the Species Analyst architecture, which uses Z39.50 protocol, an international (ISO 23950) standard defining a protocol for computer-to-computer bibliographic information retrieval from participating databases over the Internet. The first stage consists of periodically scanning through the data sources, building and updating the list of unique scientific names that are currently available for retrieval. The second stage builds a list of georeferenced records, using the list of names. The environmental layers consist of a set of geographical coverages that describe the main environmental parameters that may affect species' geographical distributions: temperature, precipitation, rainfall, solar radiation, terrain elevation and slope, and tree coverage, among others. The original climate data used by Lifemapper was obtained from the Intergovernmental Panel on Climate Change (<http://www.ipcc.ch/>) at a scale of 0.5° and processed to 1 km resolution to match the other data sets. All variables are year averages from 1961 to 1990, of the averages for the months of January and July for the same period: cloud cover, diurnal temperature range, ground-

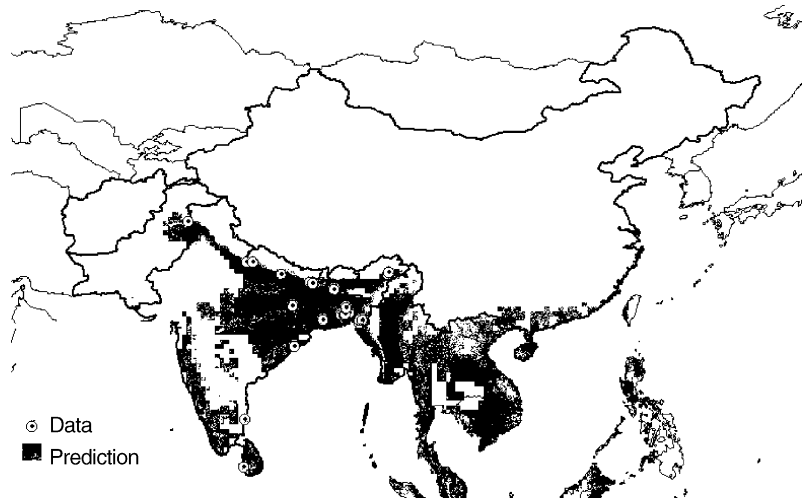


Fig. 1 – An example of Lifemapper output, the predicted distribution of a species of a fish, the spotted snakehead. Shaded areas in south-eastern Asia are where the GARP machine learning algorithm predicts favourable environmental conditions for the species given a set of observed locations (circles).

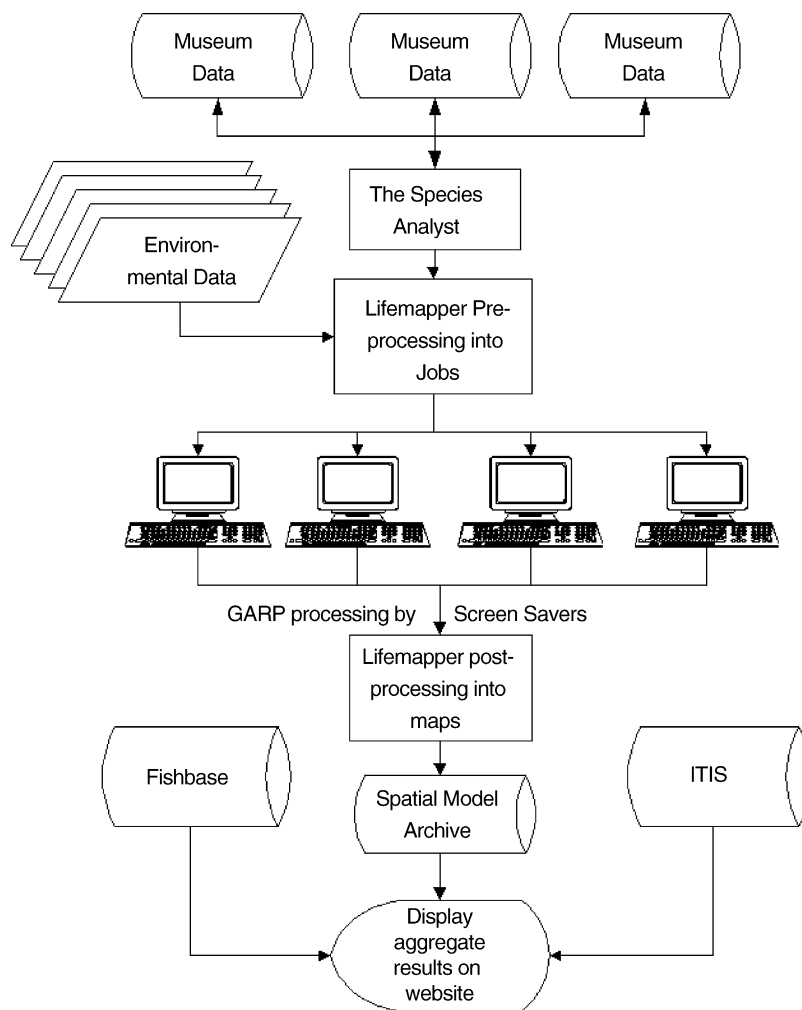


Fig. 2 – The architecture of the Lifemapper system. Pre-processing consists of data collation via distributed Internet query, processing via distributed PC screensaver clients and post-processing of results into a digital library of spatial images and other information.

frost frequency, maximum temperature, mean temperature, minimum temperature, precipitation, solar radiation, vapor pressure, wet-day frequency and winds. The terrain related layers were obtained on a scale of 1:250,000. Those layers were later processed to grids with 1 km cells.

- Aspect: describes the direction of maximum rate of change in elevation between each cell and its eight neighbors. It can essentially be thought of as the slope direction.
- Flow directions: defines the direction of flow from each cell in the digital elevation model to its steepest down-slope neighbor.
- Flow accumulation defines the amount of upstream area draining into each cell. It is essentially a measure of the upstream catchment area.
- The flow direction layer is used to define which cells flow into the target cell.
- Slope: describes the maximum change in the elevations between each cell and its eight neighbors. The slope is expressed in integer degrees of slope between 0 and 90.
- Compound topographic index or CTI, commonly referred to as the wetness index, is a function of the upstream contributing area and the slope of the landscape. The CTI is calculated using the flow accumulation (FA) layer along with the slope. In areas of no slope, a CTI value is obtained by substituting a slope of 0.001. This value is smaller than the smallest slope obtainable from a 1000 m data set with a 1 m vertical resolution.

Vegetation and land use data was obtained from University of Maryland at 1 km resolution derived from remotely sensed Landsat NDVI data: percentage of tree cover and use land cover type.

## 2.2. Processing

A genetic algorithm is a well-known machine learning method (Holland, 1975). This class of algorithms is inspired by the concept of evolution via natural selection and is based on the idea of evolving solutions to problems in a way which is analogous to the way organisms evolve. The idea is to create a set of potential solutions to a problem (the population of organisms) and then iteratively modify and test this set until a near-optimal solution is found.

In GARP, a population is a set of individual rules for predicting the presence or absence of a species at a cell. Rules are composed of chromosomes that encode the coefficients for the variables in the model. For example, each individual in a population has “chromosomes” for climate, geology, aspect and the abundance of a species. The “fitness” of each individual in a population is assessed at each iteration (“generation”) and determines the “reproductive success” of the individual. Fitness of an individual is based on the statistical significance of the rule in predicting presence and absence. A new population of candidate solutions is formed from the fittest individuals, with each individual altered from its parents by random mutation operators: point mutation and crossover. This procedure is repeated until a stopping criterion is met. GARP utilizes a rule archive that maintains a set of the best, uniquely different rules. The benefit of the archive is that the algorithm

does not converge on a single rule, or solution, but develops a set of solutions that together produce a robust predictive model.

GARP uses four types of rules simultaneously in the population of individuals: atomic rules, Bioclim rules, range rules and logit rules.

- An atomic rule uses only a single value of the variable in the precondition of the rule, e.g. “if the average annual temperature is 23 °C and the geology type is 4 then the species is present”.
- A Bioclim rule is based on the type of model used in the Bioclim program used by Nix (1986), for predicting the range of a species from their environmental tolerances. The Bioclim rule encloses the range of the environmental values of the data points in an “envelope”. The distribution of the species is predicted at those points that fall within that environmental envelope and absence predicted outside those points. For example, “If the annual average temperature is greater than 23 °C and less than or equal to 29 °C... predict presence”. The negation of a Bioclim rule can be used to predict presence or absence.
- A range rule is a generalization of the Bioclim rule. In a range rule a number of variables may be regarded as irrelevant, that is, all possible predictor variables need not be used in the rule. When applying these rules, the undefined variables are inconsequential and can have any value. Range rules also allow negation, i.e. the rule applies outside of the indicated range. For example, “If the annual average temperature is NOT greater than 23 °C and less than or equal to 29 °C... predict presence”.
- Another family of rules called logit rules is useful when species respond to the environment through environmental gradients. Logit rules are an adaptation of logistic regression models to rules. For example the logistic regression gives the output probability  $p$  that determines if a rule should be applied where  $p$  is calculated using:  $p = 1/(e^{-y} + 1)$  and  $y$  is the sum  $a + (a \times a) + b + (b \times b) + \dots + n + (n \times n)$ , where  $a \dots n$  represent variables and their coefficients. If  $p$  is greater than 0.75 then the logit rule predicts presence.

Once the set of rules is obtained, GARP evaluates each rule at each cell. Zero, one or many rules may apply to a particular cell. When many rules apply, the rule with the greatest expected accuracy determines prediction of the presence or absence. The resulting map is the prediction of the areas that would be suitable for that species. The screen saver running on a Lifemapper’s computer produces one of these maps for every “job” it computes. Many maps are produced for each particular species.

## 2.3. Post-processing

A critical requirement of any grid application based on unreliable components is to ensure the trustworthiness of the returned calculations. The system must ensure that faults in returned calculations do not affect the integrity of the overall calculation. Sources of faults come from hardware and network sources, and users might provide a system to the grid without executing the application fully.

Firstly, jobs for each species are sent out to many different processing sites. When (and if) these jobs return, measures are in place to check the data conform to the expected format, thus weeding out the dubious returns. Jobs will continue to be sent out for each species until sufficient returns are obtained.

Secondly, averaging the returned results increases the confidence in the calculation of a species distribution. Each run of GARP produces a slightly different map of the distribution of a species, due to randomness introduced in the initial and iterative stages of the genetic algorithm. This is indicative of the range of possible species distributions given a limited set of species locations and environmental variables. Averaging the results from a number of runs gives us a graduated model with values ranging from 0 to 100% that can be interpreted as probability of occurrence. So, if three maps predict presence in a particular area and one map does not, a probability of 75% is predicted at the area.

The results acquired from the screensaver processing units are stored in a database at the Lifemapper server site. These include maps of species distribution and statistics on the identity of the processing unit, the data used to develop the map and which museums contributed the data. The Lifemapper database provides a range of useful information: a resource for research scientists interested in species distributions, raw data for producing reports on the usage of contributors data and the creation of the overviews such as membership statistics. Other data sources utilized by Lifemapper at this stage include the ISIS database for standard scientific nomenclature (<http://www.sis.agr.gc.ca/pls/itisca>) and FISH-BASE (<http://www.fishbase.org>) for a listing of common names that match our scientific names.

#### 2.4. Membership management

Lifemapper is a distributed computing project, which means that the participants (Lifemappers, as we call them) contribute by downloading our screensaver, doing number crunching and enrolling new members. Members can form groups, whose members have common interests and goals. For example, a group of Kansas bird watchers can meet each other and help the Lifemapper project out at the same time by creating a group. Once the group is created, its attributes may be by species, by name or number of jobs. The website displays progress reports on numbers of data and models developed. Taxa reports show statistics on numbers of online accessible records of unique biological names. Specimen reports deal with unique specimen records that can be accessed online. These numbers vary enormously between species and may increase over time. The changes are tracked by unique taxon name additions, country from which most specimens came and museum of origin.

#### 2.5. Archive access

Users can access the archive of maps of species distributions developed by Lifemapper by searching for the species common or scientific name. A query on 'Bobcat' for example returns a listing of alternative possible species (*Lynx rufus* or *Lynx rufus rufus*), the number of contributing data sources, the number of georeferenced specimen records contributing to the model

and the number of GARP models computed. The maps can be viewed with a range of backgrounds and at a range of scales. There is provision for feedback or review of the quality of the map for each particular species. Maps can also be accessed via the Lifemapper web mapping service (WMS). Directions are given on the web site for constructing URLs that return maps as an image in GIF format to a standard web browser. The Lifemapper site also implements an open GIS consortium (OGC) compliant web service that allows formatted maps to be generated dynamically and included in web pages. A conforming OGC web service has a number of capabilities that facilitate computer-to-computer communications including standardized query strings and parameter names, and the inclusion of a function that returns the capabilities of the site in XML format metadata that allows a remote program to construct valid map requests.

### 3. Discussion

Reliability of prediction using data of variable quality is crucial to the success of analysis of ad hoc data. The data from museum databases are collected in an ad hoc manner and species are often represented by very few samples. Thus, methods are needed that perform well on unstructured surveys and small samples, are robust and efficient in use of data. Although 'robust' is not easy to define, it generally means 'perform well over a wide range of conditions'. The capacity of genetic algorithms to perform well in poorly structured domains is well known (e.g. Goldberg, 1989). The robustness of GARP has been confirmed in a number of recent studies evaluating the effects of biological data sample size and sampling bias, and variation in the types of responses of species to the environmental variables.

Due to the efforts of museums to digitize their data, the data available and numbers of species are increasing rapidly. However, given the vast number of species (millions of insects alone) and the scant data on many of them, it is imperative that the modelling method is "data efficient"—i.e. it produces accurate models using the minimum data. Data efficiency can be measured by how few data are required to develop an accurate model. Evaluations of the GARP algorithm show an average of 90% of maximum accuracy was achieved with a minimum of 10 data points (Stockwell and Peterson, 2002b). In this same study, the genetic algorithm required half as many data points as logistic linear regression and fine resolution surrogate models to achieve the same accuracy. Therefore GARP was approximately twice as data efficient at small sample sizes. It also achieved accuracy equal to or exceeding logistic regression and surrogate modelling methods on up to 100 data points. The genetic algorithm and the rule-set composed of multiple models provides an accurate 'generalized' model covering a range of data sizes.

This capacity to perform well over a range of data characteristics is also shown with respect to the environmental data. Species typically have a unimodal response to the environment, i.e. the response surface describing the suitability of habitat will have an optimal value for a particular variable such as temperature and fall to zero with increasing distance from that optimum. However, other forms of response surface

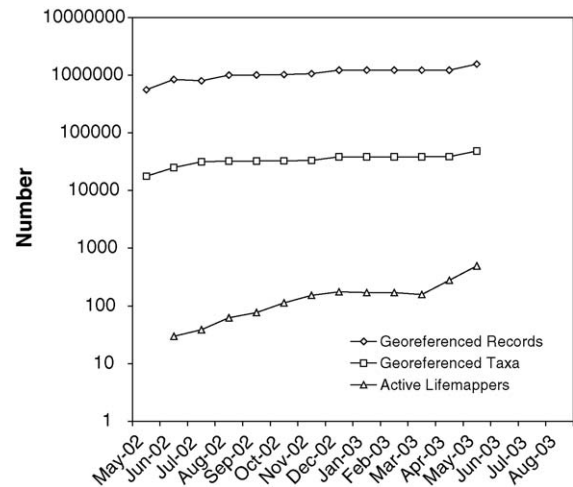
are possible. In the case of surrogate models where environmental variables are based on vegetation types, optimal suitability is related to unique values or categories in a variable. The response therefore is highly non-linear, in fact, a discontinuous or piecewise distribution. In order to make full use of the range of responses to environmental correlates a system was needed that could recognize different types of correlation that are possible with different types of environmental data. The results of Stockwell and Peterson (2002b) show the capacity of the genetic algorithm and the rule-set it develops to provide an accurate model with a range of types of environmental variables. In contrast, the accuracy of stepwise additive logistic regression models decreased with the inclusion of a variable composed of categorical vegetation types.

Finally, a pervasive problem in practical modelling is that of overfitting, or the production of models that are too specific for a given data set, and therefore perform poorly when tested on independent, new data. The iterative approach of genetic algorithms can reduce overfitting by evaluating the fitness of models repeatedly on random samples of data. The problem of overfitting is particularly severe where small data sets are not randomly selected—a situation known as bias. Again, in a comparative study using controlled amounts of bias, the GARP genetic algorithm performed better than alternative approaches (Stockwell and Peterson, 2002a).

These results provide comprehensive support for the use of the GARP algorithm in the development of this large information resource. While the results do not preclude the development of more robust methods, they show the importance of understanding the effects of data quality and evaluating robustness of prediction in alternative algorithms. Expanding the conditions under which algorithms can perform reliably is crucial for expanding the applicability of models to readily available data sources. The advantages of the genetic algorithm in GARP rely heavily on the use of machine learning methods, which along with neural nets, decision trees and adaptive agent-based learning methods are increasingly being used in ecological applications (Fielding, 1999). As a result of this work, users can have greater confidence in Lifemapper's results for application to biodiversity research, education and conservation worldwide. There are many potential uses, for example:

- Researchers will be able to model and simulate the spread of emerging diseases, plant and animal pests, or invasive species of plants and animals and their effects on natural resources, agricultural crops and human populations.
- Environmental scientists will be able to model and predict the effects of local, regional or global climate change on Earth's species of plants and animals.
- Land planners and policy makers will be able to identify the highest priority areas for biodiversity conservation.
- Teachers, students and the public will be able to discover and map their backyard biodiversity and how it might be affected by changes in rainfall or temperature or by the spread of other species.

Biodiversity science has yet to develop a complete electronic atlas of the Earth's biological diversity. Lifemapper will help science to achieve this goal, by developing an ever-increasing



**Fig. 3 – The increase in participation in Lifemapper as shown by increase in numbers of active Lifemappers, numbers of georeferenced taxa available from museums and increase in total numbers of georeferenced specimen records.**

flora and fauna of the world's taxa as data becomes available. The number of taxa in Lifemapper has been increasing explosively as more databases come online, currently greater than 100,000 species (Fig. 3). The project will assist in the use of this knowledge to better understand and conserve Earth's biological diversity, and inform environmental solutions for Earth's biological diversity.

We consider the project will provide an incentive for the museum collection community to database, clean and contribute their data to biodiversity databases such as the Species Analyst (Vieglais et al., 1998). As well as showing the scrolling display of the list of contributing museums on the screen-saver, the system also provides detailed feedback to museum data providers on how many times their data points have been viewed in maps via collection analysis statistics. More detailed feedback including possible errors in the data will be provided in the future. This relationship with the museums increases the value of participation and helps them see benefit in contributing free access to data for use by the project. The success of this strategy can be seen in the growth in number of records provided by contributing museums for use in the project, currently greater than 300,000 records (Fig. 3).

Another important benefit of Internet computing encouraging the involvement of the community in science, as can be seen by the growth in number of users contributing CPU's to the project (Fig. 3). We expect that Lifemapper will be an important tool for outreach to the general community, not only through the screen saver program, but also through a spatial query interface addressing the question "What grows/prows in my backyard/county?".

Lifemapper is an example of a real world application of genetic algorithms coupled with the computational resources of Internet computing. It will require ongoing use of large computing resources.

As such it demonstrates a dimension of problem on a par with bioinformatics, medical and other applications of large-

scale evolutionary algorithms. Lifemapper demonstrates that biodiversity researchers can cast their problems in a form suitable for execution on home computers and then persuade the public that their problem is important enough to justify the expenditure of 'free' cycles. In the future, Lifemapper will be an ongoing computational resource for addressing a wide range of computationally intensive biodiversity research projects and an indispensable component of the informatics infrastructure of biodiversity science.

---

## Acknowledgments

We thank the National Science Foundation for its support via a grant from the Division of Environmental Biology (DBI9873021).

## REFERENCES

---

- Fielding, A., 1999. *Machine Learning Methods for Ecological Applications*. Kluwer Academic Publishers.
- Goldberg, D.E., 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley.
- Holland, J.H., 1975. *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor, MI.
- Nix, H.A., 1986. A biogeographic analysis of Australian elapid snakes. In: Longmore, R. (Ed.), *Atlas of Elapid Snakes of Australia*. Australian Flora and Fauna Series Number 7. Australian Government Publishing Service, Canberra, pp. 4–15.
- Peterson, A.T., Stockwell, D.R.B., 2002. Distributional prediction based on ecological niche modelling of primary occurrence data. In: Scott, J.M., Heglund, P.J., Morrison, M., Raphael, M., Haufler, J., Wall, B. (Eds.), *Predicting Species Occurrences: Issues of Scale and Accuracy*. Island Press, Covello, CA, pp. 614–617.
- Peterson, A.T., Ortega-Heuerta, M.A., Bartley, J., Sánchez-Cordero, V., Soberón, J., Buddemeier, R.H., Stockwell, D.R.B., 2002. Future projections for Mexican faunas under global climate change scenarios. *Nature* 416, 626–629.
- Scott, M.J., Heglund, P.J., Morrison, M.L., Haufler, J.B., Raphael, M.G., Wall, W.A., Samson, F.B. (Eds.), 2002. *Predicting Species Occurrences: Issues of Scale and Accuracy*. Island Press.
- Sibley, C.G., Monroe, B.I., 1990. *Distribution and Taxonomy of Birds of the World*. Yale University Press, New Haven, Conn.
- Stockwell, D.R.B., 1999. Genetic algorithms II. In: Fielding, A.H. (Ed.), *Machine Learning Methods for Ecological Applications*. Kluwer Academic Publishers, Boston, pp. 123–144.
- Stockwell, D.R.B., Noble, I.R., 1992. Induction of sets of rules from animal distribution data: a robust and informative method of data analysis. *Math. Comput. Simulat.* 33, 385–390.
- Stockwell, D.R.B., Peters, D., 1999. The GARP modelling system: problems and solutions to automated spatial prediction. *Int. J. Geogr. Inf. Sci.* 13 (2), 143–158.
- Stockwell, D.R.B., Peterson, A.T., 2002a. Controlling bias during predictive modeling with museum data. In: Scott, J.M., Heglund, P.J., Morrison, M., Raphael, M., Haufler, J., Wall, B. (Eds.), *Predicting Species Occurrences: Issues of Scale and Accuracy*. Island Press, Covello, CA, pp. 537–546.
- Stockwell, D.R.B., Peterson, A.T., 2002b. Effects of sample size on accuracy of species distribution models. *Ecol. Model.* 148, 1–13.
- Vieglais, D.A., Stockwell, D.R.B., Cundari, C.M., Beach, J., Peterson, A.T., Krishtalka, L., 1998. The species analyst: tools enabling a comprehensive distributed biodiversity network. In: *Biodiversity, Biotechnology and Biobusiness. Proceedings of the 2nd Asia Pacific Conference on Biotechnology*, Perth, Western Australia, November 23–27.