Collective Intelligence of the Artificial Life Community on Its Own Successes, Failures, and Future

Steen Rasmussen
Self-Organizing Processes
EES-6, MS-T003
Los Alamos National Laboratory
Los Alamos, NM 87545
and
Santa Fe Institute
1399 Hyde Park Road
Santa Fe, NM 87501
steen@lanl.gov

Michael J. Raven
Reed College
3203 SE Woodstock Blvd.
Portland, OR 97202
and
Self-Organizing Processes
EES-6, MS-T003
Los Alamos National Laboratory
Los Alamos, NM 87545
ravenm@alumni.reed.edu

Gordon N. Keating
GISLab
EES-9, MS-D452
Los Alamos National Laboratory
Los Alamos, NM 87545
gkea@lanl.gov

Mark A. Bedau
Reed College
3203 SE Woodstock Blvd.
Portland, OR 97202
mab@reed.edu

Abstract We describe a novel Internet-based method for building consensus and clarifying conflicts in large stakeholder groups facing complex issues, and we use the method to survey and map the scientific and organizational perspectives of the artificial life community during the Seventh International Conference on Artificial Life (summer 2000). The issues addressed in this survey included artificial life's main successes, main failures, main open scientific questions, and main strategies for the future, as well as the benefits and pitfalls of creating a professional society for artificial life. By illuminating the artificial life community's collective perspective on these issues, this survey illustrates the value of such methods of harnessing the collective intelligence of large stakeholder groups.

Keywords
Internet, survey, consensus, consent, conflict, collective intelligence, decision support, organizational learning, professional society

1 Introduction

We describe a novel Internet-based method for building consensus and clarifying conflicts involving complex issues that large stakeholder groups must face. Then we use this method to survey and map the scientific and organizational perspectives of the artificial life community during the Seventh International Conference on Artificial Life (Artificial Life VII), held in Portland, Oregon, during August 2000. Coinciding with the advent of the new millennium, the theme of Artificial Life VII was "looking backward, looking forward," and the collective intelligence methods used here were designed to promote that theme. The survey asked two kinds of questions. One sought open-ended responses about the scientific orientation of artificial life: its main successes and failures, and its most important open scientific questions and most pressing challenges. The other questions elicited responses to a predefined list of possible functions and

© 2003 Massachusetts Institute of Technology

pitfalls of creating an international professional society for artificial life. At the conclusion of the conference, preliminary results of the survey were presented and followed by a panel discussion, a town hall meeting of all the conference delegates, and the subsequent publication of a list of fourteen important open problems in artificial life [2]. This report presents a more thorough and quantitative analysis of the results of that survey.

There are two motivations behind this report. One is to present the results of our survey, in order to convey the artificial life community's collective perspective on itself. The other motivation concerns the general method we used. We want to document and analyze this particular instance of Internet-based collective intelligence in order to facilitate the process by which groups can use information networks to create collective knowledge. We think that the power of this method is one more example of the decentralized, emergent processes illuminated by artificial life itself.

1.1 Background to the Method

The collective intelligence methods used here arise from several different traditions concerning the facilitation of citizen participation, the study of social groups, the use and development of survey methods, and investigation of the new possibilities provided by the Internet.

Citizen participation has a long tradition in public service and in large private organizations. It consists of a variety of methods for involving stakeholders in decisions that affect them. Town meetings, focus groups, and public hearings are some of the better-known methods for facilitating citizen participation. A comprehensive discussion of many classical approaches to this issue is available on the Web [3]. These traditional methods of achieving citizen participation present obvious problems for communities that are geographically distributed or that need to act quickly. The Internet shows promise for enabling both of these challenges to be addressed. Examples of some of the possibilities can be found in a discussion of a new Internet-based company-wide method for employee involvement [7].

Surveys have a long tradition as a simple quantitative method for gathering information about the lay of the land in large stakeholder groups. Traditional surveys are often formulated as specific questions with predefined possible answers (i.e., multiple choice questions). It is no surprise, then, that such survey results are very sensitive to the formulation of the questions, the available possible answers, and the context in which they appear [8]. Contacting large and widely separated survey populations and coding their responses can be difficult and tedious, so it is no surprise that there is growing interest in administering surveys over the Internet [20].

There are a variety of techniques for visualizing and analyzing information about the beliefs or opinions of a social group, organization, or community. One method is semantic networks or ontologies: ordered conceptual networks that express the shared knowledge within a group. Semantic networks represent the knowledge that enables groups to categorize, share, weigh, locate, transmit, and interpret information. Much of the knowledge in these networks is tacit, as no individual in the group is aware of its shared global structure. This tacit knowledge enables groups to solve much more complex problems than individuals could solve on their own. The study of semantic networks dates back to the first dialectic analysis in ancient Greece, and such networks form the core of most expert systems and traditional artificial intelligence applications [4, 18]. Newer applications of semantic networks include more user-friendly and effective library search systems [6]. Sociologists have studied the social implications of these and related social networks for many years (see, for example, the journal *Social Networks* and [9]). And new ways to use the Internet to examine the structure of communities are being continuously developed (see, e.g., [5, 17]).

1.2 Collective Intelligence via the Internet

The Internet is changing the way all kinds of social groups deal with the critical issues that they face. The Internet enables us to better understand the operative characteristics of communication in such groups as they reach decisions, and it enables us to develop new kinds of decision-support systems [19, 11]. The present survey is one example of this new development; it would have been much more difficult to conduct without the Internet. Distributed methods have great potential to increase the ability of organizations to achieve their goals, utilize their resources, and prepare for the future, and they easily complement our existing centralized approaches to problem solving. In the long run they may well provide new means by which society can increase its vitality and improve its quality of life.

Human society has evolved from small, separated hunting tribes to a huge, globally integrated society. This society contains a vast number of social groups, including companies, community organizations, and families (among many others). These groups have functional properties lacked by the individuals in them, and the groups are able to meet challenges that are difficult or impossible for individuals to tackle. Social groups can be found throughout the animal kingdom. The simplest explanation for why animals (ants, wolves, etc.) form groups is that this enhances the lives of the individuals. Even though we may not recognize how groups improve our lives, we still belong to them.

An increasing part of our social and intellectual dynamics is playing itself out on the Internet. The Internet is a medium of human information exchange with enormous depth and breadth. Its depth consists in capturing the full complexity of human use and creation of information and knowledge, and its breadth consists in its integration of machines, information, and people. When the full diversity of the society's intellectual dynamics is combined with the Internet's ability to quickly and accurately link information, large groups can quickly and efficiently pool their resources and coherently analyze issue complexes that were very difficult to cope with in other ways.

Traditional science and engineering approach the problem of predicting, controlling, and optimizing the behavior of a system through a premeditated linear method that typically consists of the following stages:

1. Understand how the system's behavior depends on the interconnected functioning of its components and their connections with the environment.
2. Represent the system's behavior in an approximate model (perhaps a simulation).
3. Predict the system's behavior from the model's behavior.
4. Use the model to figure out how to control those aspects of the system that bring about predefined desired goals.
5. Repeat the process, fine-tuning the model to correct for failures and improve model fidelity.

As long as we can understand how a system works, refinement of this method allows us to modify the system to meet our needs. It is desirable to follow this approach whenever possible because it enables us to modify systems quickly so as to achieve optimal system behavior.

But some systems, such as the global economy, are too complex for this traditional method. Their complexity makes their full dynamics impossible to understand, perhaps because they lack central control. In this case, the linear approach to optimization of a system is derailed at the start.

One source of inspiration for solving this problem is to take a page from nature and ask how complex distributed biological systems, such as immune systems or evolving

populations, achieve self-regulation and continual adaptation using a completely different method, one that involves no premeditation or central control. This distributed control in such systems combines emergent system dynamics with the existence at all times of a broad diversity of potential solutions. The system's overall dynamics yield the selection of candidate solutions. If no adequate solution is found within the current diversity of current candidates, the system fails and "dies." If a solution emerges, then the system dynamics will further optimize it, not by sacrificing the diversity of solutions but by changing the relative dominance of their subspecies.

Herein lies the thrust of our approach to social decision-making and control. Patterns emerge from the collective interactions of stakeholders, and these patterns can be captured easily via the Web. By making these patterns explicit and visible for the whole group, everybody gains a better understanding of the global lay of the land, and better collective decisions can be made. Since these decisions might not consciously be taken by any of the stakeholders, this method provides a foothold in the domain of unconscious decision making.

This kind of distributed emergent approach to group problem solving will never completely supplant the linear approach, because the linear approach can quickly and efficiently find optimal solutions for simpler groups. But for those complex systems that defy timely understanding, only the distributed approach can effectively search through the solution space. Distributed emergent problem solving has existed for as long as life has. By utilizing the Web, we now have the capability to understand and augment this capability in a clear and precise manner. Once the limits of the linear approach and the capability of distributed emergent solutions are known and appreciated, collective intelligence-based decision support systems should come to have a respected place in societal and scientific problem solving.

1.3 Collective Intelligence at Artificial Life VII

Artificial life as a field of study is quite young and essentially interdisciplinary. Although isolated specialists in different disciplines had (in retrospect) worked on artificial life for many decades, perhaps the first attempt to collect this disparate work and encourage its cross-pollination was the conference on "Evolution, Games, and Learning: Models for Adaptation in Machines and Nature" organized by Doyne Farmer, Alan Lapedes, Norman Packard, and Burton Wendroff in 1985 at Los Alamos. At that conference Christopher Langton coined the phrase "artificial life" [17] to describe the investigation of life-as-it-could-be, and two years later he organized the first conference called "Artificial Life." Artificial life embraces the possibility of discovering lifelike behavior in unfamiliar settings and creating new and unfamiliar forms of life, and its major aim is to understand life in all its manifestations. It goes without saying that this sweeping scope requires interdisciplinary cooperation and collaboration. The proceedings of Artificial Life VII, held in Portland, Oregon, during August 2000, includes papers by authors from biology, physics, chemistry, computer science, mathematics, psychology, economics, robotics, information science, physiology, and philosophy [1].

The interdisciplinary spread of artificial life research creates special challenges for the artificial life community. It is hard to keep abreast of relevant new work when it uses different specialized vocabularies and methodologies and is published in disparate venues, and it is hard to establish and follow high standards of scientific rigor that different disciplines with unique histories and intellectual conventions will each find acceptable. The coincidence of the Artificial Life VII conference with the birth of a new millennium provided the conference organizers with a natural opportunity to address these challenges. The organizers (Mark Bedau, John McCaskill, Norman Packard, and Steen Rasmussen) believed that artificial life would remain a vital research activity only if the community periodically looked backward and reassessed its work, so that a

foundation of recognized solid achievements could continually grow. The organizers also thought that the community must periodically look forward and identify its most important open questions, in order to promote new fruitful research avenues and track progress over time. Looking backward and forward in this way, they thought, would enable the community to renew and redefine its interdisciplinary center of gravity and to reshape the direction of future research. Hence, the theme of Artificial Life VII was "looking backward, looking forward."

Authors at Artificial Life VII were invited to relate their work explicitly to the past and future of artificial life. In addition, the conference organizers arranged to produce a coherent and precisely documented list of grand challenges in artificial life. This process started with a round-table discussion at the close of the conference, at which ten senior scientists each described one grand challenge. The elaboration of grand challenges continued in a town hall discussion in which a great number of the conference delegates participated. Subsequently, the round-table participants restructured and reorganized all of the suggestions and published a list of fourteen central open problems in artificial life [2].

The conference organizers also tried to use the Internet to promote the "looking backward, looking forward" theme. One method was to create and seed a series of threaded discussions on issues that encouraged participants to reflect broadly on artificial life's methods and accomplishments. These discussions were announced to everyone who expressed interest in Artificial Life VII, and they were linked from the conference website, but they did not spark much activity. The other method was to create a web-based survey about the science and politics of artificial life. The survey was also announced to those interested in Artificial Life VII and linked from its website. In addition, a special pitch for the survey was made at the beginning of the conference, when it was explained that the preliminary results of the survey would be presented at the close of the conference.

The round-table discussion, town hall meeting, and web-based threaded discussions and survey—all of these were interrelated means designed to probe the collective intelligence of the artificial life community about its own successes, failures, and future. This paper takes a detailed look at the results of the last of those methods: the web-based open-response survey.

2 Web-Based Conflict Clarification, Consent, and Consensus Building

In this section we describe a web-based process that can build consensus and clarify conflicts efficiently and inexpensively. By harvesting and organizing group perspectives and by feeding this information back into the group, this method harnesses the collective intelligence of the stakeholders by identifying the most important issues, showing how they are related, and showing their connection to the different stakeholder subpopulations. The method facilitates consensus building in a large stakeholder group by clarifying agreements and conflicts, and it also easily documents collective decisions.

2.1 The Basic Idea: Open-Response Survey

The method for clarifying conflict and building consensus surrounding a complex nest of issues facing a group can be broken down into the following series of steps, illustrated in Figure 1:

1. A small, diverse, and representative subset of the group designs an initial information repository, containing key questions about the problem complex on the Web. All stakeholders in the group individually review the information about the issue complex, either through the associated Web environment or through

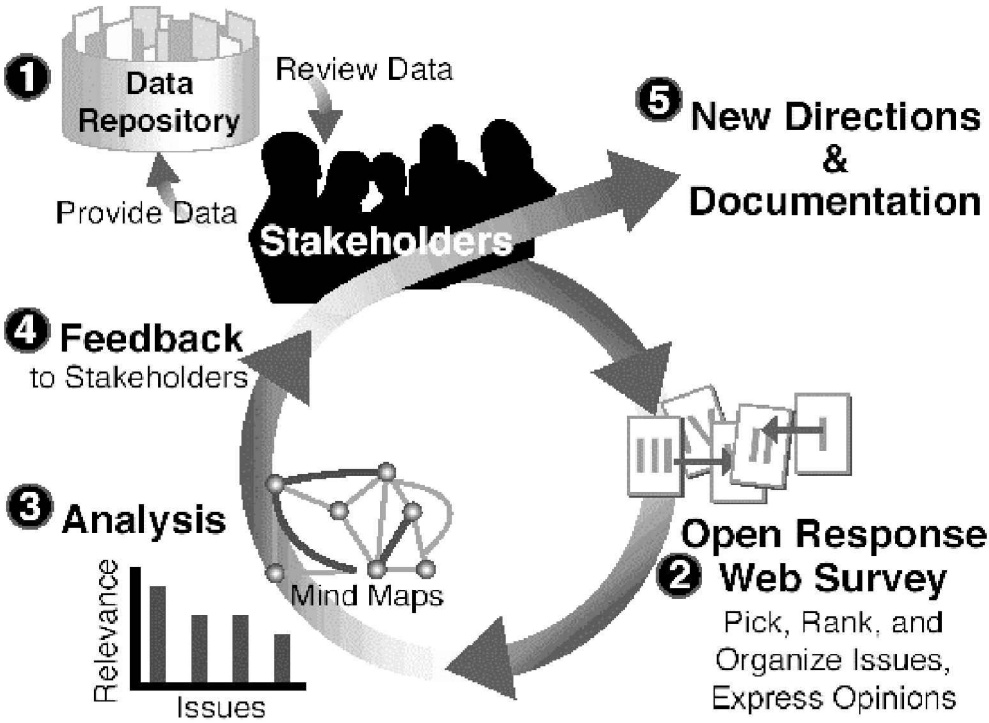

Figure 1. Schematics of the web-based consensus-building and conflict-clarification process used in connection with the Artificial Life VII Conference.

town hall meetings, media, conversations, and so on. Where possible, stakeholders add information about the problem context to the Web storehouse for review by others.

2. All stakeholders in the group rank and organize issues relevant to the problem and express their opinions about the issues through an online, open-response survey that allows freely typed input to questions. Individuals can describe new issues as well as rank those already indicated.

3. The feedback from step 2 is synthesized and analyzed to identify areas of conflict and consensus via graphical frequencies, correlations, mind maps, and other relevant plots.

4. The results of the analyses are made available to group stakeholders through the Web.

5. Plans for new directions can be developed based on the documentation of the whole consensus process.

Steps 2–4 can be repeated as the group reacts to areas of conflict and agreement, and as individuals modify their positions. (We did not do this in the present study.) Once a group has clarified its conflicts and identified its areas of consensus, it can take action on these matters. It is our contention that this sort of self-organizing collective intelligence process enables a group to make better-informed decisions about the important problem complexes that it faces.

A special aspect of the collective-intelligence process described above is the open-ended responses allowed in step 2. The open-response survey does three important things: It organizes stakeholder input along a set of broadly defined questions about the

nest of issues; it allows open-ended input; and it limits each response of each individual to a few sentences. Once the open-response data has been gathered, it is brought into a quantifiable form by coding each response into one of a finite number of response categories (here, done manually). Familiar statistical and other methods can be employed to extract information from the the open-response data once it has been categorized.

The key feature of the open-response survey is the ability to take input that is completely open in content and restricted only in length. The open-response survey can be thought of as a "fishing net" that efficiently and inexpensively catches all the worries, excitement, visions, complaints, and the like in the group and makes them available for both qualitative and quantitative analysis. This mitigates the familiar bias in traditional surveys caused by forcing all responses to be chosen from predefined answers to predefined quesitons.

2.2 Implementation

The webpages used for the open-response survey at Artificial Life VII utilized Active Server Pages technology to dynamically build HTML pages and interact with a database using JavaScript and VBScript. The open-response questions solicited feedback in the form of 255-character answers, and space is provided for several responses (usually three to five) for each question. Designing the questions to cover the problem complex logically and coherently helps impart a logical and coherent structure to the feedback. Sample answers may be provided in order to demonstrate the desired syntax or level of detail, or to stimulate thought along more than one avenue. (This was done in our survey.) This open-endedness of the feedback helps provide a natural voice to all opinions and lessens the likelihood of suppressing perspectives. This is in contrast to traditional surveys, which force participants to choose from a predefined set of answers.

Several question formats can be implemented, including agree-disagree dualities, the short-answer format described above, quantitative rankings or categories (degree of agreement or disagreement with the posed statement, etc.), and a free-response section of nearly unlimited length. We included an unlimited-space free-response section for comments outside the question structure, including reactions to the survey or the larger consensus-building process itself. It has been our experience that these free responses consistently capture important input from the stakeholders, but here we use unlimited-space free-response data only to help us interpret and clarify other non-response entries.

When all the raw survey data has been entered into a database, responses to individual questions are extracted by querying the database. These responses are categorized—in this case, manually by the authors—and then the answers in the database are recoded. The categorization process is another critical component of this method. Categorization of open-response data is a well-known challenge in the survey tradition, and the sociological survey literature is full of discussion on the issues of categorization and the validation of category sets. We do not have any special new methods for addressing this important old problem. Several of the responses in our survey could be categorized in more than one way (we return to this topic below). Nevertheless, when we performed a few random validation checks in which two different groups developed the categories independently and then compared their results, we found it surprisingly easy to reach stable categories. This holds as well for all of the other surveys that we have conducted [12–14]. So, although one should never forget that the conceptual structure extracted from the raw data is not unique, the resulting categorization still has intersubjective validity provided the categorization is consistent.

It should be noted that this method can be implemented in such a way that the responses are fully anonymous, even to those analyzing the raw data. This might be desired when sensitive political issues are being discussed, for example, and a premium is put on honest and open responses. However, if accountability and transparency are more important, it is easy to track each individual's input.

2.3 Method for Analysis of Open-Response Results

We represent formally the data for the open-response questions as follows. Let $\mathcal{N} = \{1, \ldots, N\}$ be the respondent population (stakeholders), $\mathcal{Q} = \{1, \ldots, Q\}$ be the set of distinct open-response questions on the survey, and $\mathcal{R} = \{1, \ldots, R\}$ be the set of response categories determined for a given (open-response) question. Note that the stakeholder's answers are placed in these categories after the answers are categorized as discussed in Section 2.2. Define $\Xi = (\mathbf{X}_1, \ldots, \mathbf{X}_N)$, where $\mathbf{X}_p = (\mathbf{x}_{p,1}, \ldots, \mathbf{x}_{p,Q})$ is the vector of the pth respondent's answers to the Q survey questions. Further let $\mathbf{x}_{p,q} = (x_{p,q,1}, \ldots, x_{p,q,R})$ be the vector containing the pth respondent's answers to the qth question. Thus, $x_{p,q,r}$ is the code for the number of times the pth respondent answered with category r to question q. In other words, $x_{p,q,r}$ defines the number of times the pth person picked category r under the general question q. Usually $x_{p,q,r}$ is either 0 or 1, but it may happen that a stakeholder picks closely related aspects of a single issue such that they all end up in the same category after the categorization process and thereby yield a value higher then 1.

Perhaps the simplest analysis of these data consists of extracting the frequencies of the different responses r to question q, which is given by

$$F_{\text{open}}(q, r) = \frac{\sum_p x_{p,q,r}}{\sum_p \sum_r x_{p,q,r}}$$

which, once normalized, is plotted for each of the Q main questions. Secondly, we can extract the pairwise correlations between the issues under each main question q. This correlation is given by counting all pairs (2-tuples) of issues from each individual p within the set of vectors $\boldsymbol{x}_{p,q}$. Every time person p responds with r_1 and r_2 for question q, it is counted. The counter for the connection between issues r_1 and r_2 for question q is given as

$$C_{\text{open}}(q, r_1, r_2) = \sum_p I(x_{p,q,r_1}, x_{p,q,r_2})$$

where

$$I(x_{p,q,r_1}, x_{p,q,r_2}) = \begin{cases} 1 & \text{if } x_{p,q,r_1} x_{p,q,r_2} \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

These correlations can be used to define the mind maps for the stakeholder group for a particular question. Mind maps, in particular, demonstrate not only the distribution of answers per category but also the interconnectivity of answer categories. For instance, if a respondent included answers falling into categories 1, 5, 6, and 8 in response to a given question, and if issues 1, 5, 6, and 8 are displayed as nodes in a graph, the fact that the respondent grouped issues 1, 5, 6, and 8 together can be represented by links among the nodes for those categories. Analogous information for other respondents can be pooled in the same graph, with multiple instances of nodes or links from different respondents being represented by the size of the node or link.

The full set of nodes and links for a given question constitutes a mind map for that question. This is a map not of the minds of individual respondents but of the aggregate "collective mind" (so to speak) of the group of respondents. Mind maps directly represent two kinds of consensus among the respondents: the significance of individual issues (represented by the size of a node) and the connection between issues (represented by the size of a link). Aggregating the nodes and edges into a

collective mind map leads to the emergence of larger clusters of responses that illustrate relationships among the issues raised by the group.

Our data could be given much more sophisticated statistical and graph-theoretic analysis (both off-the-shelf and custom-made), but funding limited our exploration of such methods. A more comprehensive discussion of such analysis methods is in preparation [15]. However, the first-order analyses described here illuminate significant trends in the collective perception of the artificial life community.

Finally, the diagrams resulting from the analysis may be easily used to report the areas of agreement and conflict to the stakeholder group. New positions may be formed in response to the results, and further clarification may be necessary, for example, through a new iteration of the open-response survey. The diagrams and analyses provide documentation for possible decisions based on stakeholder consensus, consent, and conflicts.

2.4 Method for Analysis of Multiple-Choice Results

We represent formally the data for the multiple-choice questions by a method that is very similar to our representation of the open-response results. Again let N be the number of respondents (stakeholders), and let Q be the number of distinct questions on the survey. The formal structure of the data is as follows. Let $\mathbf{X}_p = (x_{p,1}, \ldots, x_{p,Q})$ be the vector of the pth respondent's answers to survey questions 1 through Q. Let $x_{p,q}$ be the value of the pth respondent's answer to the qth question, where the valid responses belong to the set $\{-3, -2, -1, 0, 1, 2, 3\}$.

We present two kinds of analyses of this data. First, we plot the normalized frequency of positive responses (a response $x_{p,q}$ is positive iff $x_{p,q} > 0$) as follows:

$$F_{\text{fixed}}^{\text{pos}} = \frac{\sum_p S_{\text{pos}}(x_{p,q})}{\sum_p [S_{\text{pos}}(x_{p,q}) - S_{\text{neg}}(x_{p,q})]}$$

where

$$S_{\text{pos}}(x_{p,q}) = \begin{cases} x_{p,q} & \text{if } x_{p,q} > 0 \\ 0 & \text{otherwise} \end{cases}$$

Second, we plot the normalized frequency of negative responses (a response $x_{p,q}$ is negative iff $x_{p,q} < 0$) as follows:

$$F_{\text{fixed}}^{\text{neg}} = -\frac{\sum_i S_{\text{neg}}(x_{p,q})}{\sum_p [S_{\text{pos}}(x_{p,q}) - S_{\text{neg}}(x_{p,q})]}$$

where

$$S_{\text{neg}}(x_{p,q}) = \begin{cases} x_{p,q} & \text{if } x_{p,q} < 0 \\ 0 & \text{otherwise} \end{cases}$$

For the fixed-responses questions we plot both normalized frequencies for positive and negative responses on the same graph. Note that responses where $x_{p,q} = 0$ are not analyzed.

2.5 Problems with the Method

It is important to recognize two problems with the above-described method for building consensus and clarifying conflict. First, it is very hard to get people to give their input to any survey, traditional or open-response. The Artificial Life VII conference organizers encouraged participation by e-mailing notices before the meeting and announcing the survey several times during the conference (jokingly threatening to ban non-participants from the conference banquet). The final participation was about 30% (71 people out of about 240 conference participants), which is typical for a survey with a convenience sample (see discussion in the next section).

Second, as we mentioned above, there are inherent biases in our survey method. It is challenging to categorize the responses. In addition to taking a non-trivial amount of human time and effort, the analysts need substantial knowledge of the subject domain. Furthermore, it is impossible to eliminate all bias in the formulation of the questions and in the creation of the response categories. These biases can be mitigated if a diverse subset of the stakeholders develops the questions and checks the categories.

3 The Artificial Life VII Survey and Results

The web-based "looking backward, looking forward" survey at Artificial Life VII consisted of three sections. The first section collected demographic information about the respondent. This information included such things as contact information, organization, title, activities, age, and gender (details might vary depending on specific project requirements). The second section collected open-ended responses to questions about artificial life's main successes and failures. This part of the survey is a concrete example of the web-based open-response method described above. The third section was a traditional multiple-choice survey collecting responses about the prospect of creating a professional society for the artificial life community. A copy of the entire survey is contained in Figures 13–16, collected at the end of this paper as an appendix.

The responses to the Artificial Life VII open-response survey about scientific issues facing the artificial life community are categorized and plotted in histograms and mind maps. While the histograms illustrate the most frequently mentioned categories of answer for a given question, the mind maps provide insight into the connections among the responses. The results in this section are organized by question. We represent the results for each open-response question with both a histogram and a mind map.

Each circular node in a mind map represents a category of response to the question. The thickness of the edge of a node is proportional to the number of responses in that category, and the number within a node corresponds to the number of times this category was listed by a respondent. If someone's response to a question mentions two nodes, a line (edge) connects them. The thickness of an edge is proportional to the number of people who connected those two issues, and each edge is also labeled with the number of people who make that connection. Thus the map gives a global perspective of how the issues are related (correlated).

Since the 71 people who completed the survey represent about 30% of the population at Artificial Life VII, the survey results are meaningful for the conference population. However, we lack systematic information about who took the survey and why. Anecdotal explanations from several non-responding community members included such comments as: "I don't like filling out surveys," "I did not have time to fill it out," "Surveys are a waste of time," and "What is the purpose of this feedback?" About half of the feedback came from students and postdocs, and we would guess that the average age of the respondents was less than the average age within the artificial life community as a whole. In addition, the international artificial life community consists of many people (perhaps as many as a thousand) who did not attend Artificial Life VII.

Although roughly 5–10% of the survey responses came over the Web from people who did not attend Artificial Life VII, most artificial life stakeholders who did not attend the conference are not reflected in our survey results. So we lack firm knowledge about the biases affecting the survey results. This is a well-known feature of so-called *convenience* samples generated when survey respondents are self-selected, and it explains why convenience samples fail to provide reliable grounds for formulating statistically significant predictions and conclusions [20]. Nevertheless, they can and do suggest the lay of the land in the larger group containing the convenience sample.

In Section 3.1 we pool all survey respondents and analyze their open-ended responses about scientific issues. In Section 3.2 we analyze the open-ended responses from the two major subpopulations within the artificial life community: biologists and computer scientists. In Section 3.3 we analyze the results from all respondents about organizational issues.

3.1 Open-Response Results: Scientific Issues

Figure 2 presents the survey responses to the question "What are artificial life's most significant accomplishments?" The histogram shows that the artificial life community believes that its most important accomplishments are bottom-up modeling (labeled "Bottom-up" on the histogram and mind map, with 41 respondents) followed by sharpening the definition of life ("Def. Life," 33 respondents). Looking at the mind map, we see that these two responses are also strongly linked, with 14 respondents connecting the two responses.

Some representative examples of open-ended responses of those who thought "bottom-up modeling" was one of artificial life's most significant accomplishments are: "Showed utility of bottom-up modeling" (postdoc, complex systems), "Inspire new engineering practices" (director, computer science), "Exploring simple rules—complexity" (professor, biology), and "Models of lifelike phenomena" (student, computer science). These examples illustrate two interesting points: there is a tendency for the responders to pick formulations from the already formulated examples (first and fourth), and it is not so easy to categorize some of the responses (fourth). "Models of lifelike phenomena" could perhaps equally well be placed in a category defined as "sharpen the definition of life." However, once one has decided how to categorize a given response, it is important to make subsequent categorizations consistently. As will be clear from the mind maps for this question and in the following, the categories that are difficult to decide between will typically be highly correlated, and we see high connectivity in the mind maps.

Some representative examples of open-ended responses of those who thought "sharpen definition of life" was one of artificial life's most significant accomplishments are: "Sharpen definition of life, etc." (student, biology), "Tries to find a definition of life" (student, computer science), "Sharpen the debate about what life is" (scientist, physics), "Extraction and clarification of fundamental principles of life" (student, biology). Taken together, these two categories of main accomplishments may be interpreted as agreement on the broader question of investigating "life as it could be" using simple, local rules in computer simulations.

The next most important accomplishments of artificial life, according to the survey, are the development of a better understanding of evolution ("Evolution," 28 responses), the development of successful interdisciplinary collaborations ("Interdisciplinarity," 27 responses), a better understanding of emergent phenomena ("Emergence," 27 responses), and, lagging a little behind, attempting to understand life by creating it ("Create," 23 responses). Although the mind map connects almost every issue to every other issue, these six most popular responses to the first question define the dominating component in the mind map. This high degree of connectivity may also

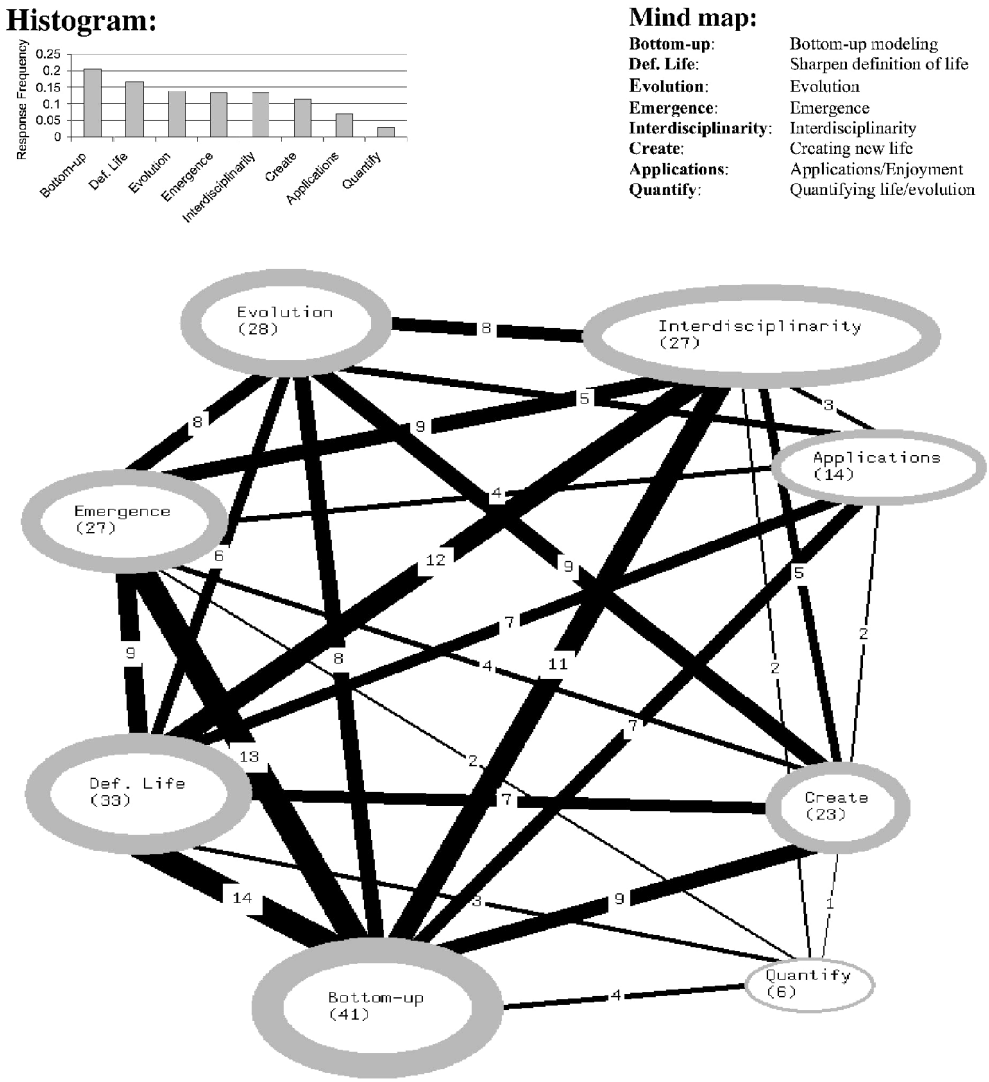

Figure 2. Responses to the question: What are artificial life's most significant accomplishments?

be interpreted as a reflection of the highly interdisciplinary character of the research activities in artificial life.

The last two responses reflected significant accomplishments involving a variety of specific engineering or entertainment applications ("Applications," 14 responses) and the development of quantitative measures for comparing natural and artificial system adaptation and evolution ("Quantify," 6 responses).

Figure 3 depicts responses to the question "What are artificial life's most significant failures?" The histogram shows that the three most frequent responses were too little theoretical or experimental grounding for the work done ("No rigor," 45 responses), no coherent agreement on which scientific problems the community should address ("No direction," 39 responses), and insufficient connection to other scientific fields ("Unrelated," 31 responses). The mind map shows a striking triangle of consensus between these three response categories (19, 15, and 15 people connecting them).

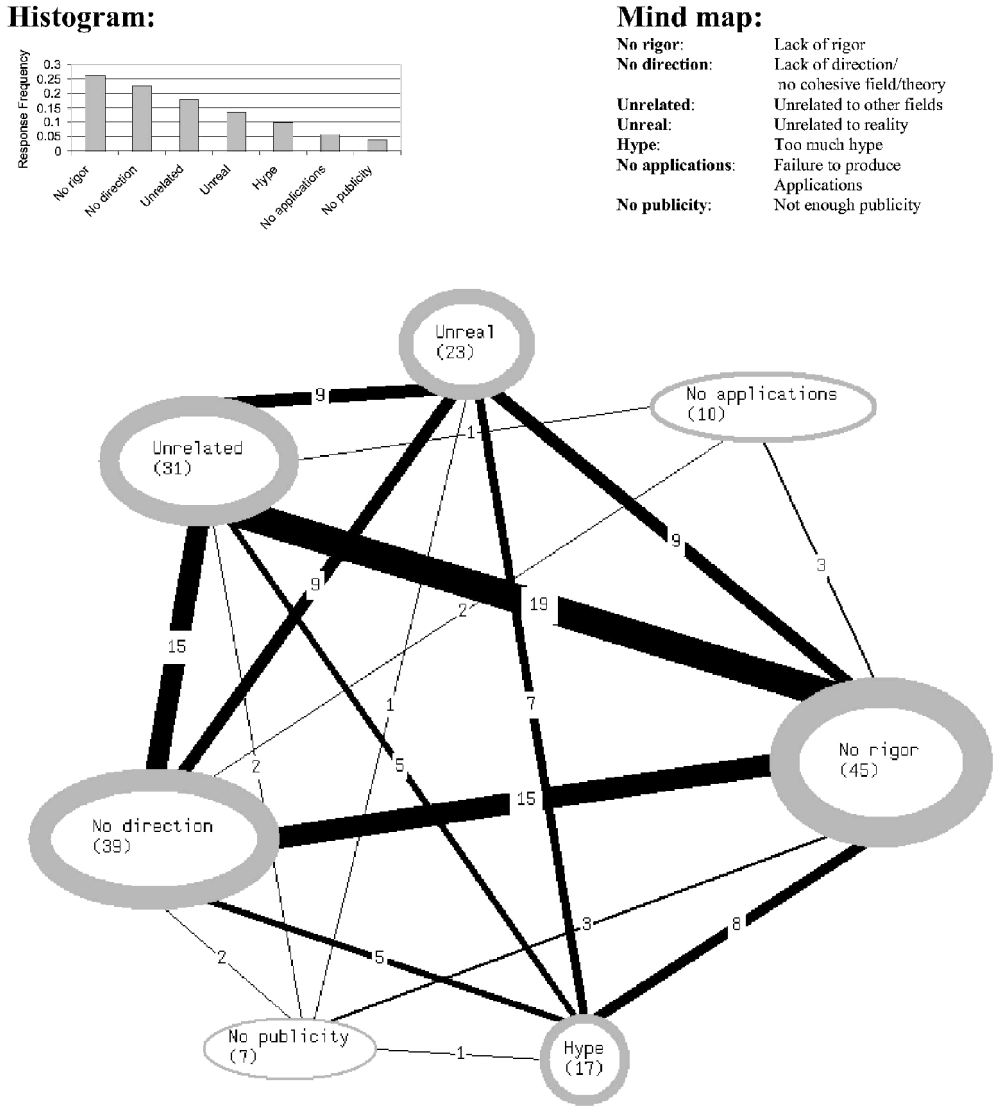

Figure 3. Responses to the question: What are artificial life's most significant failures?

With such a strong degree of consensus it should be easy to address these community problems (see also the results in Figure 5).

Some other failures identified by the community, in descending order, are that artificial life research is insufficiently related to reality ("Unreal," 23 responses), that it is associated with too much hype ("Hype," 17 responses), that no "killer" applications have been developed within the field ("Applications," 10 responses), and that not enough good publicity has come out about the community ("No publicity," 7 responses).

The responses to the question "Which scientific issues or questions is it most important for artificial life to address in the future?" are shown in Figure 4. Respondents provided a wide range of answers, which perhaps is not surprising, since the question concerns the uncertain future rather than the much-discussed past. We needed to

Histogram:

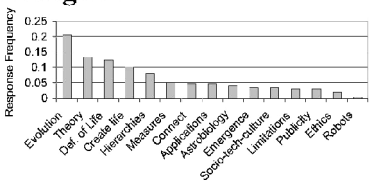

Mind map:

Evolution:	Open-ended evolution
Theory:	Theoretical foundations
Def. of Life:	Definition of life
Create Life:	Create life
Hierarchies:	Dynamical Hierarchies
Measures:	Quantitative measures
Connect:	Connect models with reality
Applications:	Applications
Astrobiology:	Extraterrestrial life/astrobiology
Emergence:	Emergence
Socio-tech-culture:	Socio-tech-culture
Limitations:	Overcome computational limitations
Publicity:	More public awareness/publications
Ethics:	Ethics/morality
Robots:	Robots

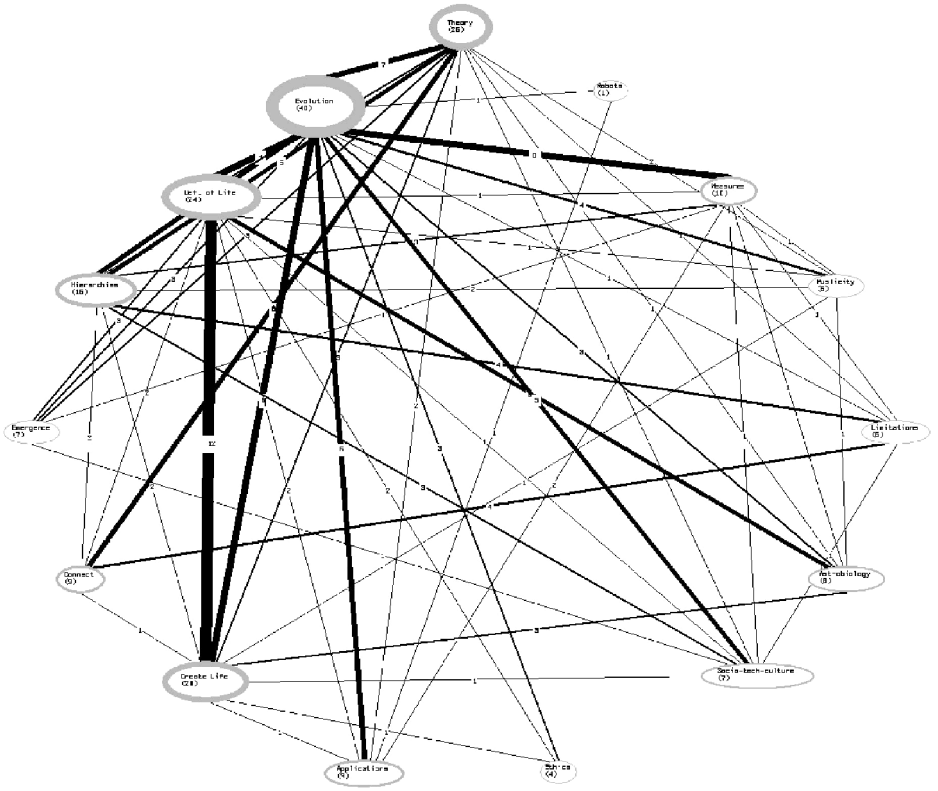

Figure 4. Responses to the question: What scientific issues is it most important for artificial life to address? The large size of the map makes the text in the nodes very small. See text for details.

create fifteen response categories to capture the diversity of feedback to this question. Nevertheless, the responses still show a significant structure.

Looking first to the histogram, we can see three groups of response frequencies. The first group consists of the single most frequent category of response: those who thought open-ended evolution was the key issue to address ("Evolution," 48 responses). The second group consists of four responses: developing more theory to supplant the current excessive ad hoc work in the community ("Theory," 28 responses), developing a deeper understanding of life either through devising a better definition of life ("Definition of life," 24 responses) or by creating life ("Create life," 20 responses), and better understanding dynamical hierarchies ("Dynamical hierarchies," 18 responses). The third group consists of the remaining, least frequent response categories.

The mind map consists of many mostly weak connections between response categories. Since this question had a larger number of response categories, one would expect more connections of lesser strength. Indeed, only a single connection achieved a strength greater than 8, namely the connection between defining life and creating it. Similarity of subject matter easily explains the close relationship between these distinct topics.

The large group of low frequency responses includes quantitative measures of natural and artificial evolution and adaptation ("Quantitative," 10 responses), one or more important applications ("Applications," 9 responses), contributing to astrobiology ("Astrobiology," 8 responses), clarifying the notion of emergence ("Emergence," 7 responses), contributing to a better understanding of socio-technical and cultural phenomena ("Socio-tech-culture," 7 responses), surmounting the computational barriers of detailed bottom-up simulations ("Limitations," 6 responses), ensuring and coordinating good publicity ("Publicity," 6 responses), addressing the social, ethical, and religious issues associated with artificial life activities ("Ethics," 4 responses), and finally, making intelligent humanlike "robots" ("Robots," 1 response).

Summarizing the broad pattern in responses to this question, understanding open-ended evolution stands out as a clear grand scientific challenge. The next most important challenges are strengthening our theoretical foundation (recall that theory was identified as a significant weakness in Figure 3) and better understanding life and dynamical hierarchies. Since most of these issues were not strongly connected, we might conclude that the community believes that they can be addressed relatively independently.

Figure 5 shows the responses to the question "What should artificial life do to address its most important scientific issues?" Increasing interdisciplinary collaboration is by far the most frequent recommendation ("Collaboration," 47 responses). The next most frequent recommendations are to focus attention and effort on the scientifically most important issues ("Focus," 33 responses; recall the previous two questions), to find and help create more funding for artificial life research ("Funding," 23 responses), to preserve the community's historical "big tent" at meetings and the journal, welcoming and valuing scientific diversity ("Big tent," 16 responses), and to provide more significant theoretical and experimental foundations for research ("Theory," 13 responses).

One can draw a number of tentative conclusions about artificial life's future from these results. First, the mind map reinforces the central importance of involving theoretical and experimental collaborators from the established fields (chemistry, biology, physics, etc.), because this is mainly what connects the other three dominant recommendations: scientific focus, external funding, and a big tent (although "Funding" and "Focus" are also rather strongly connected, as one could expect). Second, focusing on the key scientific issues is important to prevent diffuse, ill-defined, or frivolous activities from draining momentum. The paper on key open problems [2] is one constructive step toward sharpening this focus. Third, artificial life faces significant challenges when

Histogram and legend:

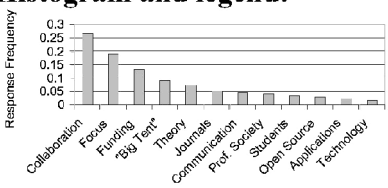

Mind Map:

Collaboration:	Interdisciplinary collaboration
Focus:	Focus on key questions
Funding:	More funding
"Big Tent":	"Big Tent" approach
Theory:	More theory (not ad hoc)
Journals:	Better/more journals
Communication:	More communication/discussion
Prof. Society:	Professional society
Students:	Student involvement
Open Source:	Open source/standardization
Applications:	More applications
Technology:	Better technology

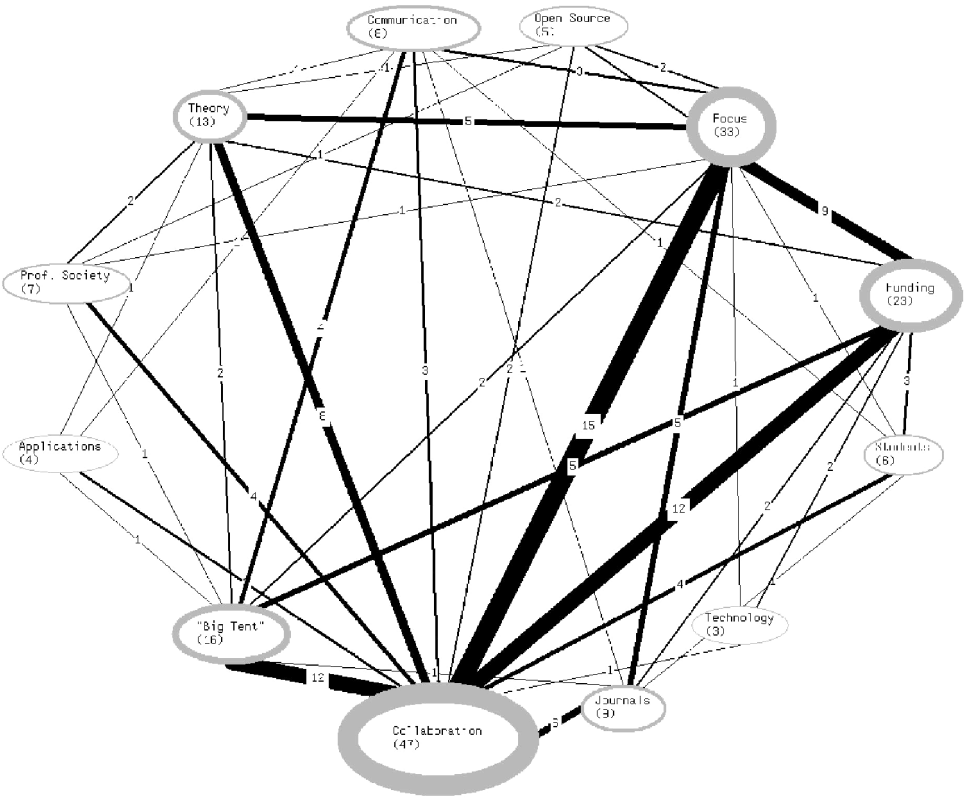

Figure 5. Responses to the question: How can artificial life solve its most significant scientific issues?

seeking research funding, since its activities are interdisciplinary but still loosely connected to traditional fields. So we need to develop strategies for making funding agencies aware of our strengths, perhaps by means of the newly established International Society for Artificial Life, and also by calling attention to the past successes that have spun off from artificial life and founded narrower, more easily identifiable and fundable research communities (evolutionary computing, genetic programming, etc.). Finally, as we become more rigorous and focused, we need to encourage cross-disciplinary activities and reach out to new scientific partners.

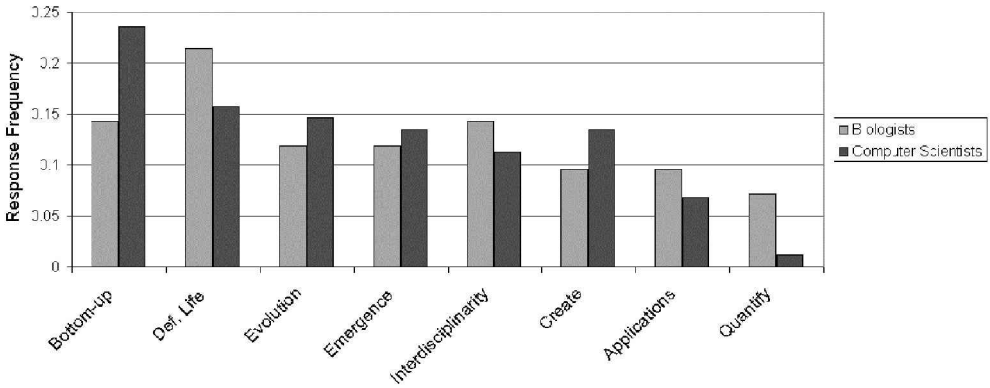

Figure 6. Comparison of how biologists and computer scientists responded to the question: What are artificial life's most significant accomplishments?

3.2 Open-Response Results: Scientific Issues for Biologists versus Computer Scientists

Among the classes of fields to which respondents belonged, biology and computer science were the largest: out of a total of 71 respondents, 35 said they were computer scientists and 16 said they were biologists. Despite the relatively small number of respondents, it is interesting and instructive to compare responses from biologists and computer scientists. Below we compare their responses to the four open-ended response questions (discussed above).

The most salient difference between biologists and computer scientists on artificial life's most significant accomplishment is that they switch the top accomplishment (Figure 6). The biologists stress sharpening the definition of life, while computer scientists think bottom-up modeling is the most important accomplishment. The frequencies for the remaining response categories were roughly the same for both populations. Nevertheless, the mind maps reveal subtle but significant differences between them (Figure 10). Biologists most strongly connect the definition of life with bottom-up models, interdisciplinarity, and applications. Computer scientists most strongly connect bottom-up modeling with emergence and interdisciplinarity.

Biologists and computer scientists also switch the rank of artificial life's most significant failure. Biologists stress lack of rigor, while computer scientists emphasize lack of direction (Figure 7). Biologists also believe that insulation from other well-established fields ("Unrelated") was a significant failure, while computer scientists judged this a relatively minor failure. Finally, computer scientists viewed the lack of significant applications as a minor failure, while biologists did not consider it a failure at all. Again the mind map (Figure 10) reveals another interesting difference between the two populations. The biologists' mind map is dominated by a single strong connection between the lack of rigor and the lack of relation to traditional fields. The computer scientists' mind map has an equally strong connection between the lack of rigor and the lack of direction.

Biologists overwhelmingly view open-ended evolution to be artificial life's most significant scientific challenge (Figure 8), with creating life a distant second. Computer scientists also think open-ended evolution is the most significant scientific issue, but this is followed closely by the definition of life, dynamical hierarchies, and stronger theoretical foundations. Also note that the computer scientists' list of significant issues is larger than the biologists' list, probably partly due to their larger population size. Larger population size is also reflected in the computer scientists' mind map being "noisier" (Figure 10).

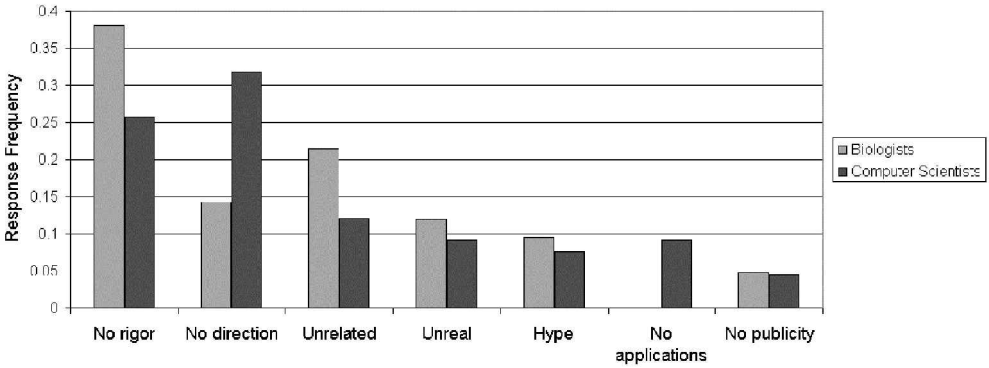

Figure 7. Comparison of how biologists and computer scientists responded to the question: What are artificial life's most significant failures?

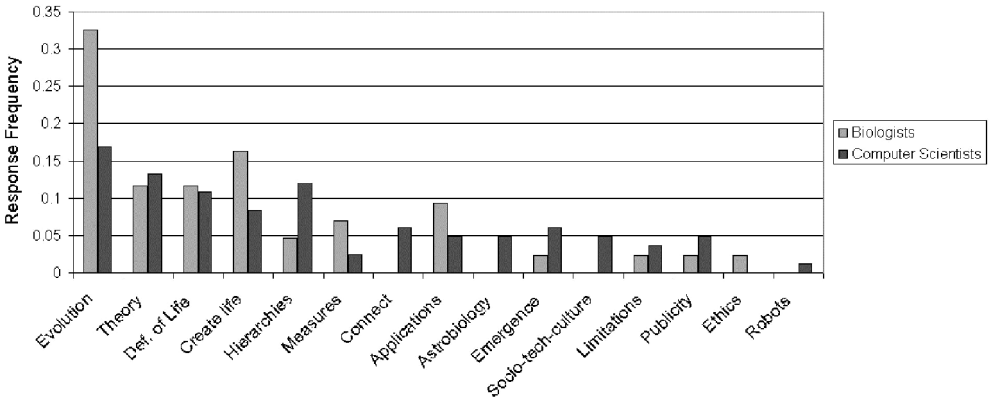

Figure 8. Comparison of how biologists and computer scientists responded to the question: What scientific issues is it most important for artificial life to address?

The largest difference between the biologists and computer scientists shows up in their recommendations for how to make artificial life thrive. Biologists overwhelmingly think that focus on the key issues is crucial, while computer scientists overwhelmingly choose interdisciplinary collaboration (Figure 9). The mind map indicates further significant differences (Figure 10). Biologists connect focus with funding and with the journal *Artificial Life*. This might be interpreted as the suggestion that the journal should foster a focus on the important issues, in order to increase research funding. Biologists do not view interdisciplinary collaboration as crucial, perhaps because they view themselves as experts on life and how to investigate it. The mind map for computer scientists has interdisciplinary collaboration at its center, connecting the four next most important issues. One might venture the interpretation that collaboration is important for computer scientists because they recognize that, while they are not experts on life, they have a methodology that can shed important light on life when properly deployed.

3.3 Results on Organizational Issues

The maturation of new scientific communities often involves the formation of a professional society. Discussion about such a society had taken place at a number of the conferences on artificial life preceding Artificial Life VII, and marked divisions had

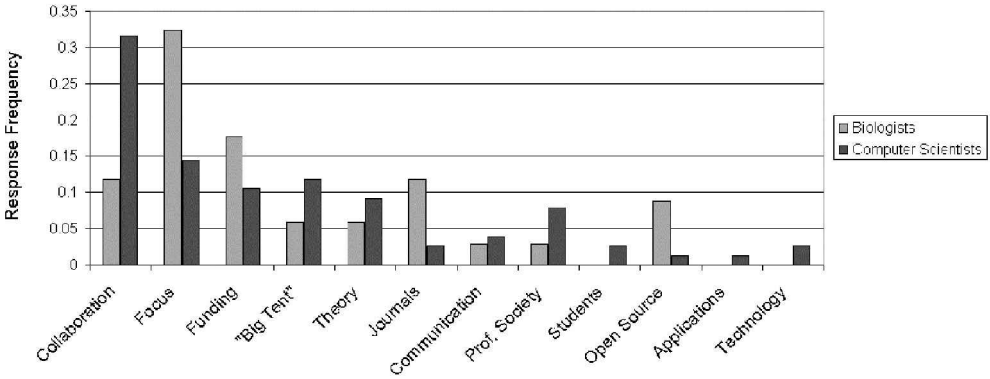

Figure 9. Comparison of how biologists and computer scientists responded to the question: How can artificial life solve its most significant scientific issues?

surfaced at the previous year's European Conference on Artificial Life (ECAL'99) in Lausanne, Switzerland. Some people felt quite strongly that the future success of artificial life hinged on a professional society. Others felt equally strongly that the problems a society would cure were imaginary or exaggerated, and that a society would generate serious new problems. So the organizers of the ALife VII survey decided to air and measure the nature and extent of the sentiment for and against a professional society for the artificial life community.

This part of the survey consisted of two multiple-choice questions. The first asked respondents to rank the desirability (from "most desired" to "least desired") of a number of possible functions of a professional society (another option was "no opinion"). The following are the thirteen possible functions of the society that were listed under this question:

1. A society could oversee the editorship of the journal *Artificial Life*.
2. A society could automate the process of subscribing to *Artificial Life*.
3. A society could provide discounted registration at artificial life conferences.
4. A society could determine the location and timing of the artificial life conferences.
5. A society could reduce the duplication of organizational support for artificial life conferences.
6. A society could provide startup funds for artificial life conferences.
7. A society could advertise conferences and other events of interest to the artificial life community.
8. A society could oversee and manage a website for the artificial life community.
9. A society could reduce the hype surrounding artificial life and promote high scientific standards for research.
10. A society could encourage support for artificial life research by governmental sources as well as private businesses.
11. A society could establish links with related professional societies.
12. A society could help students find postdocs and jobs in academia and business.

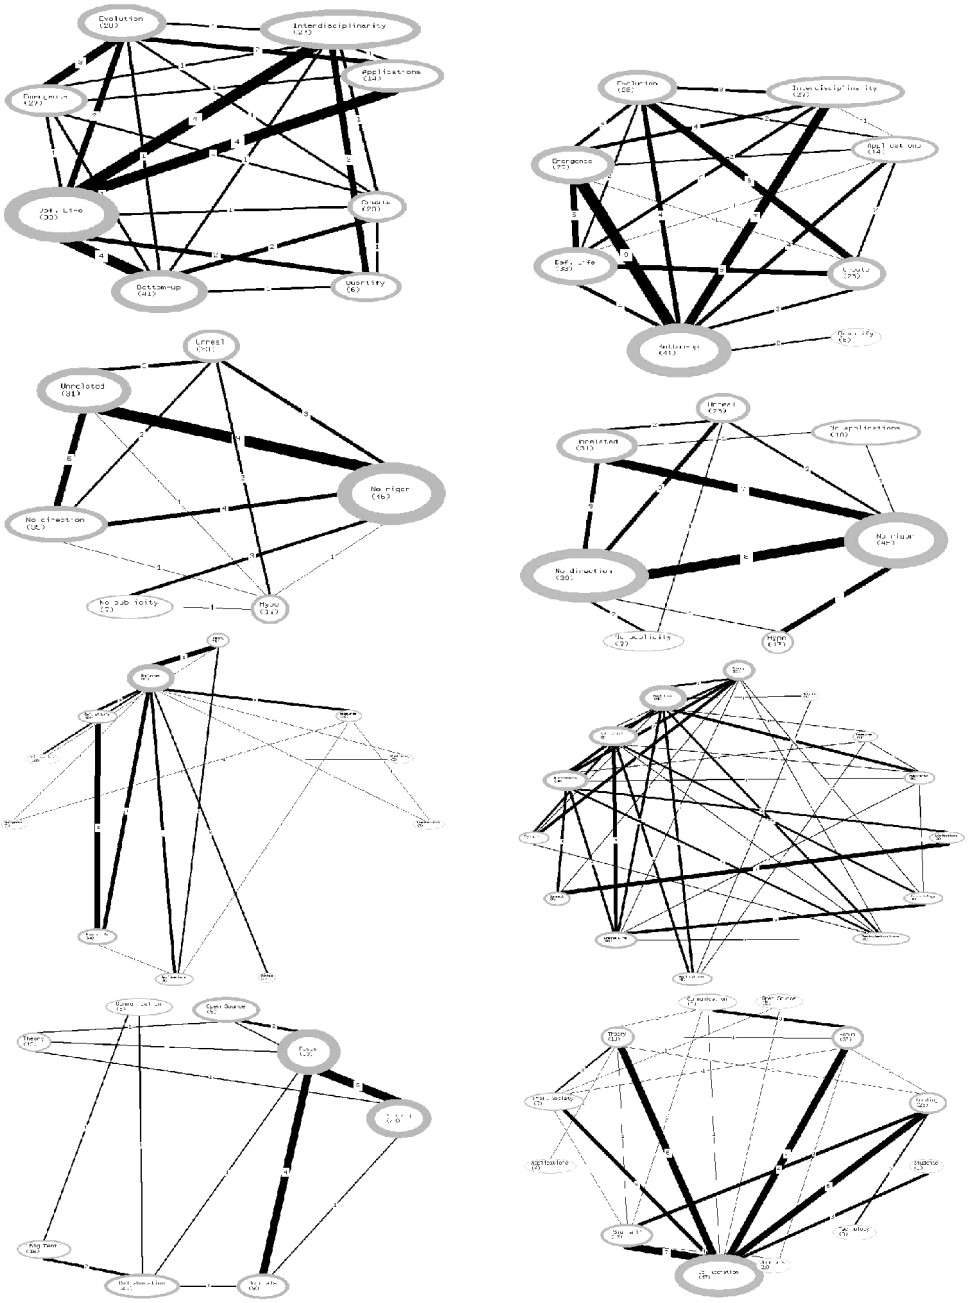

Figure 10. Comparison of mind maps for biologists (left) and computer scientists (right). The small size makes the text in the nodes very small, but the important comparison is between the global structure in the mind maps. The arrangement of nodes in each pair of biology-versus-computer-science mind maps is the same, and it exactly corresponds to the arrangement of nodes in the corresponding mind maps in Figures 2–5. The mind map on the top row corresponds to Figure 2, the next row to Figure 3, the next row to Figure 4, and the bottom row to Figure 5.

13. A society could add professional legitimacy to artificial life, especially for those with professional positions in other fields.

This list of possible functions was intended to include the main purported benefits of a society emphasized in the artificial life community's previous discussions on the topic.

The second multiple-choice question asked respondents to rank the seriousness (from "huge problem" to "no problem") of a list of possible problems that might be created by a professional artificial life society (another option was "no opinion"). As with the previous question, the options listed were intended to reflect the main worries that had emerged in previous discussions. The following are the six possible problems that were listed under this question:

1. A society would perpetuate the field after it is no longer useful.
2. The personal agendas of those running the society would dominate the field.
3. A society would not respect or preserve the scientific diversity that has historically characterized the artificial life community.
4. A society would smother the creativity and openness provided by artificial life's big tent.
5. Nobody (or nobody who is qualified) will want to do the work necessary to run the society.
6. The artificial life community does not need any more organization ("If it's not broken, don't fix it").

Figure 11 summarizes the responses concerning the possible advantages of a professional society. The histogram suggests two main conclusions. First, there are many perceived advantages of the society. Reducing hype is ranked highest, but it surpasses providing professional legitimacy only slightly, and the frequency of the remaining eleven possible advantages fall off slowly and gradually, with discounted registration at artificial life conferences having the lowest score. But even discounted registration is viewed as quite desirable. Second, there is a strong degree of consensus about the possible advantages of a society. Disagreement is reflected as significant values both above and below the 0 line in the histogram. There is almost no sense that any of the thirteen possible functions were undesirable.

Figure 12 summarizes the responses concerning the problems that a professional society might create. Three main conclusions stand out from these results. The first concerns the artificial life community's worries about a professional society. Most worrisome is the potential dominance of the personal agendas of those running the society. Next most worrisome is that unqualified people would run the society. And each of the other possible worries listed in the survey bothered a non-trivial number of respondents. The second conclusion concerns what the community on balance does *not* find worrisome (indicated by the bar below the 0 line in the histogram). In particular, the balance of the community does not worry so much that a professional society would stifle creativity, or that it would perpetuate the field beyond its usefulness, or that it would dampen the community's scientific diversity. The third conclusion concerns the degree of consensus about the pitfalls of a professional society. The virtual unanimity about the benefits of a society contrast with significant disagreement about the problems it might generate, with dampened diversity and dampened creativity generating the most disagreement. Perhaps this lack of consensus explains the heated disagreement about a professional society at ECAL'99.

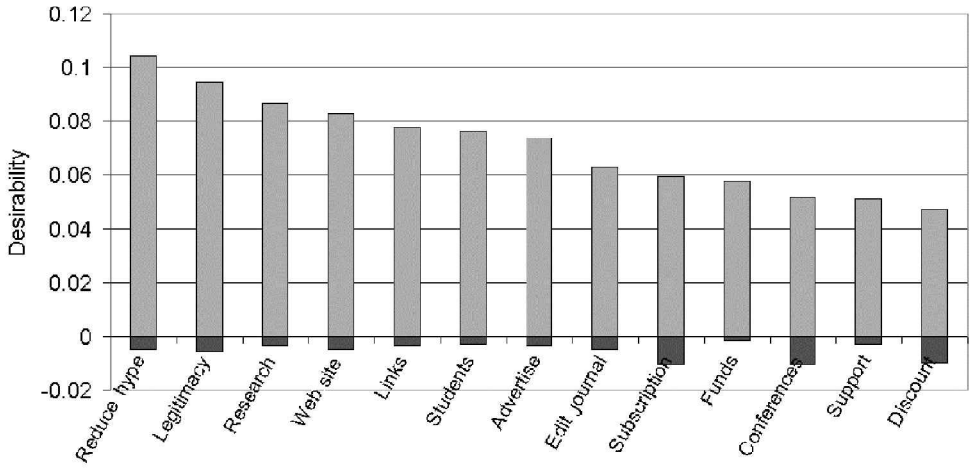

Legend

Reduce hype :	Reduce hype and promote high standards for research.
Legitimacy:	Add professional legitimacy to field, especially for those researchers grounded in other fields.
Research:	Encourage industrial and government support for research.
Web site	Oversee and manage a central web site.
Links:	Establish links with other professional societies.
Students:	Help students find postdocs and jobs in academia and business.
Advertise:	Advertise conferences and events of interest.
Edit journal :	Oversee the editorship of the journal.
Subscription:	Automate subscription to the journal.
Funds:	Provide startup funds for conferences.
Conferences:	Determine location and timing of conferences.
Support:	Reduce duplication of conference organizational support.
Discount:	Discount registration at conferences.

Figure 11. Responses on the desirability of possible functions of a professional artificial life society. Above the 0 line are shown those who thought a particular function was a desirable feature of a society. Below the 0 line are shown those who thought the function was an undesirable feature of a society.

4 Discussion and Conclusion

We draw two kinds of conclusions from the Artificial Life VII survey. One kind concerns what we can learn about web-based collective intelligence methods for clarifying conflicts and building consensus. The other kind concerns what we can learn about artificial life's successes, failures, and future. We will discuss each in turn.

It is well known that groups with more than about a dozen members have a difficult time functioning smoothly. How, then, are large and widely distributed stakeholder communities to govern themselves and adapt effectively? Web-based collective intelligence methods can be part of the answer. The results of the Artificial Life VII survey demonstrate how such methods can help large, complex stakeholder groups to clarify conflicts and build consensus, mostly by providing a rather detailed and quantitative "lay of the land." The Web made it easy to aggregate an open-ended variety of stakeholder sentiments on key issues, and familiar data analysis tools such as mind maps made it easy to visualize and analyze those aggregate opinions. Mirroring the results back to the stakeholder group (as in this paper) closes the loop and enables group opinion itself to reshape, adapt, and evolve as the group learns the dimensions of its own harmony and discord.

The present experiment only hints at the full potential of these methods. These methods could be applied to many other stakeholder groups. The methods could also

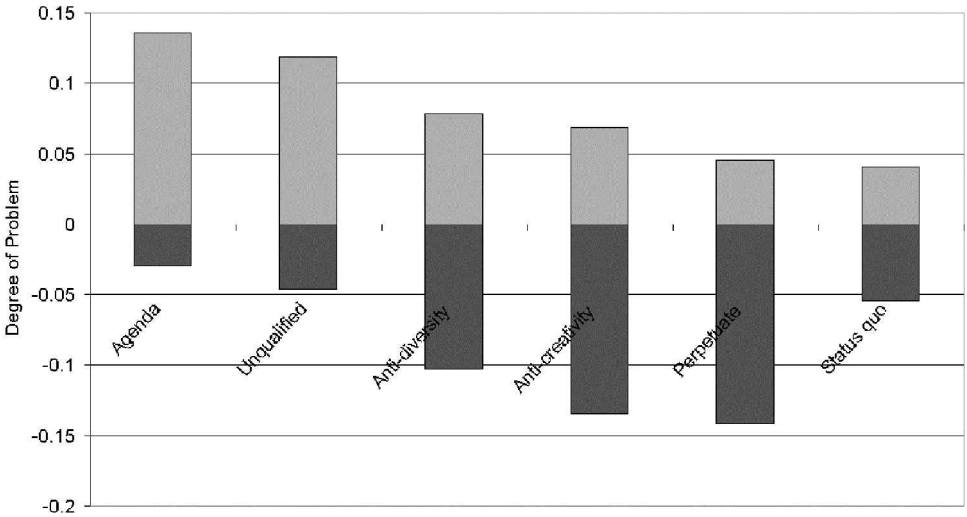

Legend

Agenda:	The personal agendas of those running the society would be too dominant.
Unqualified:	Nobody (or nobody qualified) will want to do all the necessary work.
Anti-diversity:	The society would not respect scientific diversity.
Anti-creativity:	The society would smother the creativity and openness that make it so attractive.
Perpetuate:	A professional society would perpetuate the field beyond its usefulness.
Status quo:	Things are fine now; we don't need more organization.

Figure 12. Responses to the request to indicate the seriousness of possible problems created by a professional artificial life society. Above the 0 line are shown those who thought a particular problem was serious. Below the 0 line are shown those who thought the was not serious.

be extended and generalized in a number of ways, including the following:

1. The present study involved just one circuit of the collective intelligence process (recall Figure 1), with the publication of this paper closing the loop by presenting the results back to the stakeholder group. But the process could be iterated. For example, after publishing these results, we could use them to craft a revised open-ended survey with new questions that build on key lessons from the first survey (e.g., questions about how to implement the key recommendations that emerged). This iterated application of open-ended collective methods could help the group adapt quickly and intelligently.

2. The present study has just scratched the surface in analyzing the collected data; much more could be done. In our survey it would be interesting, for example, to compare the responses of those new to artificial life and those who have been involved from the 1980s. This is one area where further iterations of our survey might be very illuminating. Furthermore, the data could be analyzed across the different questions, and clusters in the artificial life community could be identified. There are many more sophisticated methods of statistical analysis, some of them ready for immediate use off the shelf, and others that could be developed specifically for the case at hand. But application of more sophisticated analyses might need to await the collection of new data drawn from an unbiased random sample of respondents.

3. One of the main bottlenecks in the present survey process is determining how to categorize open-ended responses and subsequently code them. A contextually defined free-text language-parsing algorithm would enable this front end of the survey to be automated, at least partially, and this would vastly increase the

possible size of the stakeholder groups that could participate in the process. It is probably desirable to have humans at least double-check the categories used to code results. Random samples picked from the very large automatically categorized results would probably identify most potential biases and errors introduced by automated response categorization.

4. Currently all graphing and statistical analysis are done manually, but they could also be automated. Automation would facilitate the rapid graphical representation of survey results online, and eventually this analysis could even be done in real time, for example with search capabilities to display different certain subpopulations and the like.

Implementing proposals 3 and 4 would vastly multiply the power of these collective intelligence methods. If the input were automatically categorized, analyzed, and displayed, one would immediately be rewarded with the group's collective knowledge upon completing the survey. Upon completing the survey, a stakeholder could immediately see how her or his data fits into the global picture. We expect that this way of automating the process and making it more dynamic would significantly increase and enhance stakeholder participation. It would also facilitate the iterative process suggested in proposal 1.

Turning to the specific results of the present survey, it is not surprising that the artificial life community's biologists and computer scientists differ on the key scientific issues facing artificial life. But they largely agree about the community's main successes, its main failures, and where its future lies. There is a high degree of consensus in the community as a whole about the main problems of the past. Fundamentally these consist of three interrelated issues: insufficient scientific rigor, insufficient coherence of scientific attention, and insufficient integration with traditional disciplines.

Artificial life's two main successes are viewed as the advent and development of bottom-up modeling and sharpening the definition of life, followed by advancing our understanding of the dynamics of evolution and of emergent phenomena and fostering interdisciplinary research. The scientific issue thought most critically to need attention, by an overwhelming margin, is open-ended evolution, followed by more attention to theory and to synthesizing life and dynamical hierarchies. It is striking that each of these issues was also highlighted in the paper on key open problems [2]. The four main ways seen to address these key issues successfully are collaborations with traditional disciplines, focusing on key scientific questions, developing more research funding, and nurturing our big tent culture.

Finally, the whole community sees many advantages in a professional artificial life society, which can help coordinate our activities and thus enable us to better address the issues above.

The community has already made progress on many of the issues highlighted by this survey. For one thing, the key open problems paper [2] is already helping the community focus on core scientific issues. This effect could be seen at Artificial Life VIII in Sydney, Australia. Many of papers at that conference explicitly connected their contribution to those grand challenges, and the conference included a round-table discussion assessing progress on those challenges. In addition, the International Society for Artificial Life (ISAL) has now been created. Its membership is growing, as is the range of functions it is providing for the artificial life community. These functions include overseeing the journal *Artificial Life*, overseeing the artificial life conferences, and creating and maintaining an ISAL website (http://www.alife.org). ISAL is also helping to build connections with other conferences in which artificial life figures, with new efforts to link to traditional conferences of related traditional disciplines (such as evo-

lutionary biology and ecology). Though probably most of the artificial life community hopes for more from ISAL, its efforts to date are clearly a significant step in a positive direction. And the interest in joining ISAL's Board of Directors expressed by over a dozen people at Artificial Life VIII is also promising. The key question now is how ISAL can best focus the enthusiasm of the artificial life community.

Acknowledgments

This project grew out of work under the Urban Security Initiative [10] at Los Alamos, and it was developed with the help of Johan Bollen, Claes Andersson, Sudha Maheshwari, and Andrias O'Callagham. Helen Coculelis, Noah Goldstein, and Norman Johnson provided valuable discussion in earlier similar projects. At the Artificial Life VII meeting in Portland, Oregon, John McCaskill helped us categorize the preliminary survey results. Norman Packard helped us design the original Web instrument for polling the artificial life community. Martin Nilsson has given us feedback on future mathematical analysis of the categorized data. John Pock helped us understand the pitfalls of surveys, especially those allowing open-ended responses, and he pointed out useful literature on Internet surveys. Jenny Cocq has been engaged in the progress from the beginning, Pierre Dotson has given concrete ideas to improve the feedback fraction, and Norman Packard has suggested ways to make a simple natural language parser. Johan Bollen and Clif Bowen provided comments on an earlier draft. We thank all these people for their generosity and helpful insights.

References

1. Bedau, M. A., McCaskill, J. S., Packard, N. H., & Rasmussen, S. (Eds.) (2000). *Artificial Life VII: Proceedings of the Seventh International Conference on the Simulation and Synthesis of Living Systems.* Cambridge, MA: MIT Press.

2. Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., Harvey, I., Ikegami, T., Kaneko, K., & Ray, T. S. (2000). Open problems in artificial life. *Artificial Life 6*, 363–376.

3. Bleiker, H., & Bleiker, A. (2000). *Citizen participation handbook.* http://www.ipmp-bleiker.com.

4. Bobrow, D. G. (Ed.) (1994). *Artificial Intelligence in Perspective.* Cambridge, MA: MIT Press.

5. Bollen, J., & Heilighen, F. (1998). A system to restructure hypertext networks into valid user models. *The New Review of Hypermedia and Multimedia 4*, 189–213.

6. Bollen, J., & Luce, R. (2002). Evaluation of digital library impact and user communities by analysis of usage patterns. *D-Lib Magazine 8*(6), June.

7. Feder, B. (2000). IBM meets with 52,600 virtually. http://www.nyt.com/2001/05/28/technology/28BLUE.html.

8. Fowler, F. J. (1993). *Survey Research Methods.* Newbury Park, CA: Sage Publications.

9. Freeman, L. C. (2000). Visualizing social networks. *Journal of Social Structure* 1. http://www2.heinz.cmu.edu/project/INSNA/joss/vsn.html.

10. Heiken, G., Valentine, G., Brown, M., Rasmussen, S., George, D., Green, R., Jones, E., Olsen, K., & Andersson, C. (2000). Modeling cities: The Los Alamos urban security initiative. *Public Works Management and Policy, 4*, 198–212.

11. Johnson, N., Rasmussen, S., Joslyn, C., Rocha, L., Smith, S., & Kantor, M. (1998). Symbiotic intelligence: Knowledge on distributed networks driven by human interactions. In C. Adami, R. Belew, H. Kitano, & C. Taylor (Eds.), *Artificial Life VI* (pp. 403–407). Cambridge, MA: MIT Press.

12. Keating, G., Rasmussen, S., & Raven, M. (2001). *Consensus building tools for GIS design* (Report LA-13894-MS). Los Alamos National Laboratory.

13. Keating, G., Rasmussen, S., & Raven, M. (2001). Web-based consensus building and conflict clarification for EES division's strategic planning process: New technical directions (Report LA-UR-02-3830). Los Alamos National Laboratory.

14. Keating, G., Rasmussen, S., Raven, M., Tso, E., Cocq, J., & Dotson, P. (2001). Use of web-based consensus building and conflict clarification process for the Navajo Nation Governmental Efficiency Study, an appendix to *ETD environmental consulting, report of the Navajo Nation Council Evaluation*, submitted to the Office of Navajo Government Development, Window Rock, AZ (Report LA-UR-01-6207). Los Alamos National Laboratory.

15. Keating, G., Rasmussen, S., Raven, M. J., Nilsson, M., & Bedau, M. A. (2003). Consensus building and conflict clarification processes for collective decision support. In preparation.

16. Langton, C. G. (1986). Studying artificial life with cellular automata. *Physica D, 22*, 120–149.

17. Marsden, P., & Bollen, J. (1998). Help advertising evolve: Clone consumer thought patterns. *Ad Map 34*, 37–39.

18. Pereira, F. C. N., & Grosz, B. J. (Eds.) (1994). *Natural language processing*. Cambridge, MA: MIT Press.

19. Rasmussen, S., & Johnson, N. (1998). Self-organization at and around the Internet (Report LA-UR-98-2549). Los Alamos National Laboratory.

20. Schonlau, M., Fricker, R. D., Jr., & Elliott, M. N. (2001). *Conducting research surveys via e-mail and the Web*. RAND News and Publications.

Appendix

The entire survey is shown in Figures 13–16.

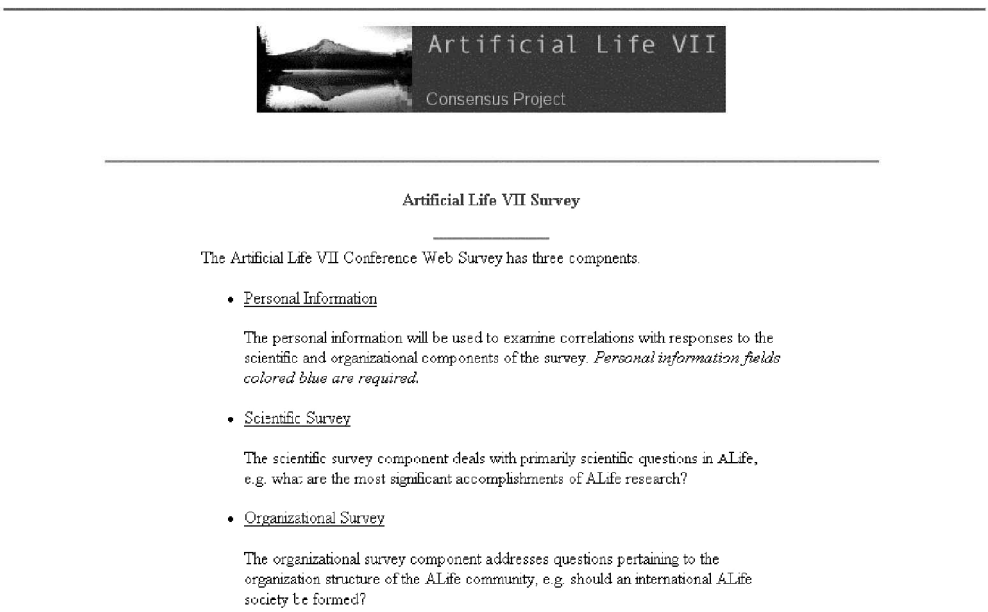

Figure 13. The first page of the Artificial Life VII survey, giving a brief introduction to the three parts.

Personal Information

Personal Web Page:	
Contact Address:	
Country:	(please select)
Name of Organization:	
Type of Organization:	(please select) If your organization type is not listed, please specify:
Size of Organization:	(please select)
Position at Organization:	
Main Field of Research/Employment:	(please select) If your field is not listed, please specify:
How long have you been doing ALife?:	(please select)
Organization's Web Page:	
Personal Interests/Research	
Research Interests of your Organization	

Figure 14. The demographic information collected in the survey.

ALife7 Scientific Survey

What are ALife's three most significant accomplishments? (type in your answer)

Accomplishment #1:

Accomplishment #2:

Accomplishment #3:

Use the following sample responses as suggestions for your own, if you wish. You do not have to use them, however, and are free to respond any way you choose.

- Generative complexity from simple rules
- Simple models of life-like phenomena
- Experimental exploration of evolving systems
- Quantitative comparison of artificial and natural system evolution
- Popularized bottom-up models in other fields
- Sharpened debate on definition of life, etc.
- Fostered interdisciplinary synthetic study of life-like systems

Where are ALife's three most significant failures? (type in your answer)

Failure #1:

Failure #2:

Failure #3:

Use the following sample responses as suggestions for your own, if you wish. You do not have to use them, however, and are free to respond any way you choose.

- Failure to engage related fields (e.g., Biology, Chemistry) enough
- No clear and cohesive definition of the field
- Too little attention to real living systems
- Too much hype
- Too little comparison between artificial and natural systems

What are the three most significant issues or problems which will be most important to ALife now and in the forseeable future? (type in your answer)

Current Issue #1:

Current Issue #2:

Current Issue #3:

Use the following sample responses as suggestions for your own, if you wish. You do not have to use them, however, and are free to respond any way you choose.

- Definition of life (universal requirements of life)
- Practical model of open-ended evolution
- Practical model of successive/hierarchical emergence (dynamical hierarchies)
- Definition of statistics to "see" open-ended evolution
- To create life in vitro
- To develop a unifying theory/framework/simulation of life
- Can life exist without genetics?
- How could/would life look at other planets/solar systems?

What are the three most significant things ALife will need to do to solve these issues/problems? (type in your answer)

Need #1:

Need #2:

Need #3:

Use the following sample responses as suggestions for your own, if you wish. You do not have to use them, however, and are free to respond any way you choose.

- Create research collaborations/conferences including researchers/theorists/experimentalists
- Create funding lines for ALife research through NSF/NIF/NASA/DOE/...etc...
- Create a business network to help support ALife efforts
- Write a special issue in the ALife journal on ALife applications and successes
- Be selective and only focus on a few key issues
- Keep a broad all inclusive approach

Please use the space below for free response.

Figure 15. The open-response component of the survey, which gathered information about artificial life's scientific successes, failures, open problems, and strategies for future success.

ALife7 Organizational Survey

Rank the desirability of the following possible functions of a professional artificial life society.	Most desired			Neutral			Least desired	No opinion
Oversee the editorship of the Artificial Life journal.	○	○	○	○	○	○	○	○
Automatic subscription to the Artificial Life journal.	○	○	○	○	○	○	○	○
Discounted registration at artificial life conferences.	○	○	○	○	○	○	○	○
Determine location and timing of artificial life conferences.	○	○	○	○	○	○	○	○
Reduce duplication of conference organizational support.	○	○	○	○	○	○	○	○
Provide start-up funds for artificial life conferences.	○	○	○	○	○	○	○	○
Advertize conferences and events of interest.	○	○	○	○	○	○	○	○
Oversee and manage a central artificial life web site.	○	○	○	○	○	○	○	○
Reduce hype and promote high standards for research.	○	○	○	○	○	○	○	○
Encourage industrial and government support for research.	○	○	○	○	○	○	○	○
Establish links with other professional societies.	○	○	○	○	○	○	○	○
Help students find post-docs and jobs in academia and business.	○	○	○	○	○	○	○	○
Add professional legitimacy to artificial life, especially for those researchers grounded in other fields.	○	○	○	○	○	○	○	○
	Most desired			Neutral			Least desired	No opinion

Rank the seriousness of the following possible problems that might be created by a professional artificial life society.	Huge problem						No problem	No opinion
A professional society would perpetuate the field beyond its usefulness	○	○	○	○	○	○	○	○
The personal agendas of those running the society would be too dominant.	○	○	○	○	○	○	○	○
The society would not respect scientific diversity.	○	○	○	○	○	○	○	○
The society will smother the creativity and openness that makes artificial life so attractive.	○	○	○	○	○	○	○	○
Nobody (or nobody qualified) will want to do all the necessary work	○	○	○	○	○	○	○	○
Things are fine now, we don't need more organization.	○	○	○	○	○	○	○	○
	Huge problem						No problem	No opinion

What direction should future ALife research take?

Please use the space below for free response.

Submit Reset

Figure 16. The predefined-option (multiple choice) component of the survey, which gathered information about the pros and cons of a professional artificial life society.