

**Graph-Theoretic Measure for Active iGAs:
Interaction Sizing and Parallel Evaluation
Ensemble**

**Xavier Llorà, Noriko Imafuji Yasui, and
David E. Goldberg
IlligAL Report No. 2008003
March, 2008**

Illinois Genetic Algorithms Laboratory
University of Illinois at Urbana-Champaign
117 Transportation Building
104 S. Mathews Avenue
Urbana, IL 61801
Office: (217) 333-2346
Fax: (217) 244-5705

Graph-Theoretic Measure for Active iGAs: Interaction Sizing and Parallel Evaluation Ensemble

Xavier Llorà¹, Noriko Imafuji Yasui², and David E. Goldberg²

¹National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
1205 W. Clark Street, Urbana, IL 61801
xllora@illigal.ge.uiuc.edu

²Illinois Genetic Algorithms Laboratory
University of Illinois at Urbana-Champaign
104 S. Mathews Ave. Urbana, IL 61801
{kumara,tianliyu,deg}@illigal.ge.uiuc.edu

March 10, 2008

Abstract

Since their inception, active interactive genetic algorithms have successfully combat user evaluation fatigue induced by repetitive evaluation. Their success originates on building models of the user preferences based on partial-order graphs to create a numeric synthetic fitness. Active interactive genetic algorithms can easily reduce up to seven times the number of evaluations required from the user by optimizing such a synthetic fitness. However, despite basic understanding of the underlying mechanisms, active interactive genetic algorithms still lack of principled understanding of what properties make a partial ordering graph a successful model of user preferences. Also, there has been little research conducted about how to integrate together the contribution of different users to successfully capitalize on parallelized evaluation schemes. This paper addresses both issues describing (1) what properties make a partial-order graph a success, and (2) how partial-order graphs obtained from different users can be merged meaningfully.

1 Introduction

Traditional interactive genetic algorithms (iGA) (Caldwell and Johnston, 1991) require the user to provide a large number of evaluations to achieve good quality solutions. Due to the repetitiveness of the tasks, interactive sessions may last for hours usually inducing fatigue and frustration making harder for the user to maintain the evaluation criteria during the whole evolutionary process. A key contribution of active interactive genetic algorithms (aiGAs) (Llorà et al., 2005) is their capacity to build models of user preferences to generate *educated guesses*—promising solutions. Presenting *educated guesses* for user evaluation helps reducing the interaction fatigue—the user can obtain high-quality solutions faster—and the frustration—generating high-quality solutions allows aiGA to avoid the repetitive display of poor solutions that may discourage the user. Such an approach has shown that user evaluation can be reduced the number of user evaluation from 3 to 7 times.

A key element to the success of the aiGA paradigm, as shown in Llorà et al. (2005), is its ability to reconstruct a global ordering out of the partial-order graph built using the user evaluations. However, there is little knowledge about what make a partial-order graph $G = \langle V, E \rangle$ a good candidate to generate an accurate enough global ordering, or ranking, of the solutions contained in V . Moreover, when facing the design of aiGA there are three variables that need to be taken

into account: the length of the solutions (ℓ), the growth ratio of nodes ($\Delta(V)$), and the growth ratio of edges ($\Delta(E)$). Usually, given a certain problem to solve ℓ is fixed, but we still need to decide how often we need to generate a new node or show a new comparison to the user. In this paper we present a first analytic study based on graph-theoretic measures to separate accurate partial-ordering graphs (the ones that can provide an accurate ranking that will lead to accurate solutions) from the inaccurate ones. As a result, we have identified a lower bound based on the density of the partial-order graph. If the partial-order graph density is below that boundary, the algorithm will not be able to provide high quality solutions—as empirically shown in Llorà et al. (2006).

Another daunting challenge for any interactive genetic algorithm is how to integrate the results obtained from different runs. If you are trying to solve a problem by means of an iGA, you may end having multiple run results that you need to combine. Moreover, those results may also be coming from different users and, thus, it is highly likely that results may differ—or even contradict. aiGAs can help mitigate this problem by assembling multiple partial-order graphs together. Contradictory evaluations may be introduced by *user mistakes*, *perceptually indistinguishable solutions*, *contradictory targets*, and *criteria shifts*. In this paper we start addressing the first category of such assembling contradictions: user mistakes. We show that such mistakes can be probabilistically modeled as binomial distributions and, thus, simple contradiction resolution policies can be put in place to maximize the quality of the resulting ensemble.

The rest of this paper is organized as follows. Section 2 presents relevant background about active interactive genetic algorithms. It provides an overview of its mechanisms and how they relate to the work presented in this paper. Then, section 3 presents a first analysis of the partial-order graphs. It also shows a lower bound for the partial-order graph density that needs to be satisfied to be able to compute an accurate synthetic fitness. We also present in section 4 a first approach to assembling partial-order graph obtained from different interactive session to create a final solution for the problem addressed using active interactive genetic algorithms. Finally, section 5 presents the conclusions for the work presented in this paper.

2 Active Interactive Genetic Algorithms

Dawkin’s Blind Watchmaker program (Dawkins, 1986) and the Faceprints system developed at New Mexico State University (Caldwell and Johnston, 1991) are two early examples of iGAs. For example, in Faceprints, the system replaces the role of a human sketch artist in evolving the faces of criminal suspects from witness recollection. Faces are encoded as binary strings where subcodes represent different facial features (nose type, mouth type, hair type, etc.). Each full chromosome maps to a face and the population of chromosome is presented to the human critic who is asked to determine how close the face resembles that of the criminal. This subjective ten-point scale is used to drive the evolution of subsequent generations of faces, and in a relatively short time, the GA arrives at a reasonable facsimile of the correct face.

The use of interactive genetic algorithms allow the fusion of human and computer efforts for problem solving (Takagi, 2001). However, putting the evaluation process into the hands of a user sets up a different scenario when compared to normal optimization. Takagi (2001) presented a review of research efforts related to the iGAs challenges. These research areas included: (1) discrete fitness value input method, (2) prediction of fitness values, (3) interface for dynamic tasks, (4) acceleration of iGAs convergence, (5) combination of evolutionary and non-evolutionary computation, (6) active intervention, and (7) theoretical research. A detail description of these topics

is beyond the scope of this paper and detailed descriptions can be found elsewhere (Takagi, 2001; Llorà et al., 2005).

Unlike in traditional evolutionary algorithms with objective fitness measures, one of the daunting challenges of iGAs is effective methods of combating user fatigue. Even for moderately-sized problems, iGAs may require a few hundred to a few thousand fitness evaluations, which is highly improbable—oftentimes even impossible—for users to evaluate. This places a premium on a variety of *efficiency-enhancement techniques* (Goldberg, 2002), particularly *evaluation relaxation* (Sastry, 2001). In evaluation relaxation schemes, the computationally costly, but accurate function evaluation is replaced by a cheap, but less accurate surrogate function. A serious stumbling block in developing surrogate fitness functions in iGAs is the absence of computable fitness function. Additionally, the user evaluation is relative and user preference might change over time. Hence, existing evaluation-relaxation methods fall short and cannot effectively model user fitness function.

2.1 Key Components

Active interactive genetic algorithms (aiGAs) (Llorà et al., 2005) proposed to model user preferences in order to create a relaxation scheme. aiGAs propose a method which consists of three major components:

1. *Partial ordering*: The qualitative decisions made by the user about relative solution quality is used to generate partial ordering of solutions.
2. *Induced complete order*: The concepts of non-domination and domination count from multi-objective evolutionary algorithms (Goldberg, 1989) to induce a complete order of the solutions in the population based on their partial ordering.
3. *Surrogate function via support vector machine*: The induced order is used to assign ranks to the solutions and use them in a support vector machine (SVM) to create a surrogate fitness function that effectively models user fitness.

The key element of an aiGA is its synthetic fitness function. The minimal scenario for collecting meaningful domain-independent information from the user is provided by a binary tournament scheme ($s = 2$) (Goldberg et al., 1989). User evaluations introduce a partial order among the solutions presented so far—in this paper, the synthetic representation of the weights configurations. This partial order can be made explicit by using a partial-ordering graph ensemble $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ —as firstly suggested in Llorà et al. (2005). A vertex in \mathcal{V} represents the solutions presented to the user, whereas the edges in \mathcal{E} represent the partial-ordering evaluations provided by the user. Given two solutions $\{s_1, s_2\} \in \mathcal{V}$ the user is able to provide three possible outcomes: *i*) $s_1 > s_2$, *ii*) $s_1 < s_2$, and *iii*) $s_1 = s_2$ —or *equal/don't know/don't care*. Such a graph \mathcal{G} can be transformed into a normalized graph \mathcal{G}' containing only *bigger than* relations. A global ordering measure may be computed using a heuristic based on two dominance measures, δ and ϕ , inspired by multiobjective optimization (Coello-Coello, 1998; Deb et al., 2000). Let $\delta(v)$ be the number of different nodes present on the paths departing from vertex v . Analogously, $\phi(v)$ is defined as the number of different nodes present on the paths arriving to v . The estimated fitness of a given solution v may be computed as $\hat{f}(v) = \delta(v) - \phi(v)$. Intuitively, the more solutions a vertex v dominates (is *greater than*), the greater the fitness. Otherwise, the more solutions dominate (are *greater than*) a solution v , the smaller the fitness. The final global estimated ranking $\hat{r}(v)$ is obtained sorting the vertex $v \in \mathcal{V}$ by $\hat{f}(v)$. This global estimate is used to train a ε -SVM for creating the synthetic surrogate fitness (Llorà et al., 2005). By optimizing such a synthetic fitness we can obtain *educated guesses* about

Table 1: Algorithmic description of the aiGA model proposed by Llorà et al (2005).

-
1. Create an empty directed graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$.
 2. Create 2^h random initial solutions (\mathcal{V} set).
 3. Create the hierarchical tournament set \mathcal{T} using the available solutions in \mathcal{V} .
 4. Present the tournaments in \mathcal{T} to the user and update the partialordering in \mathcal{E} .
 5. Estimate $\hat{r}(v)$ for each $v \in \mathcal{V}$.
 6. Train the surrogate ε -SVM surrogate synthetic fitness based on \mathcal{G} and $\hat{r}(v)$.
 7. Optimize the ε -SVM synthetic fitness using the cGA.
 8. Create a \mathcal{S}' set with 2^{h-1} new different solutions, where $\mathcal{V} \cap \mathcal{V}' = \emptyset$, sampling out of the probabilistic model evolved by cGA.
 9. Create hierarchical tournament set \mathcal{T}' with $2^h - 1$ tournaments using 2^{h-1} solutions in \mathcal{S} and 2^{h-1} solutions in \mathcal{V}' .
 10. $\mathcal{V} \leftarrow \mathcal{V} \cup \mathcal{V}'$
 11. $\mathcal{T} \leftarrow \mathcal{T} \cup \mathcal{T}'$
 12. Go to 4 while not converged.
-

the user preferences. The optimization step is conducted by using the compact genetic algorithm (cGA) (Harik et al., 1998) obtaining a probabilistic model of the user preferences.

2.2 Interesting properties

Llorà et al. (2005) pointed out that partial-ordering graphs have an interesting property: given a normalized partial-ordering graph \mathcal{G}' , if a vertex v appears more than once in a path of $\delta(v)$ or $\phi(v)$, then a cycle in such graph exists, i.e., it represents an inconsistency in the user evaluations. Thus, due to the *greater than* relations, the consistency of the user evaluations can be identified. This property is the basis of the consistency metric. In order to compute such a measure the authors need two components: *i*) cycle detection capabilities for a given graph \mathcal{G}' at time t (\mathcal{G}'^t), and *ii*) an heuristic to quantify how much inconsistency the detected cycle is introducing. A detailed explanation of this property can be found elsewhere (Llorà et al., 2006).

The estimated fitness is a key component of aiGA. It assigns a numeric estimate to each solution presented to the user and, hence, it allows the creation of synthetic fitness surrogates that allow estimating the quality of unevaluated—a key component to the probabilistic model building component mentioned below. It also allows measuring the error of the synthetic fitness surrogate when compared to the original estimated values. Such error measuring acts as a sanity check of the surrogate, since an erroneous surrogate will lead to unreliable models of the user preferences.

The original aiGA proposed by Llorà et al. (2005) used the cGA (Harik et al., 1998) to obtain a probabilistic model of the user preferences by optimizing the surrogate fitness function constructed. This process is repeated every generation—as shown in table 1. Hence, during interaction between the user and the aiGA several probabilistic models are obtained. The original aiGA discards the previous model when a new one is obtained. However, the sequence of models provides interesting insides about the evolution of the user preferences along the run. For instance, sustained criteria on the evaluation process by the user will be reflected by slowly changing sequences of models. However, drastic changes on the user criteria may show up as sequences of highly different model sequences.




Figure 1: Hierarchical tournament selection scheme used by aiGA. The parameter h defines the height of the tournament tree. The first 2^{h-1} solutions are drawn at random without replacement from the set of already available solutions. Solution 2^h is the best ranked solution so far, Finally, the rest of the solutions are promising solutions obtained after optimizing the synthetic fitness built on the previous iteration.

2.3 Problematic Incremental Usage

The original aiGA model proposed by Llorà et al. (2005) assumes a generational approach to the interactive process—as show in table 1. Later on, Llorà et al. (2006) proposed an incremental approach to the collection of user evaluations. Their approach was based on maintaining the steady-state approach of the original aiGA, at each evaluation step, a tournament is performed between two previously evaluated solutions that were not compared directly, or between a previously evaluated solution and a solution sampled from the probabilistic model of the user preferences.

Such a tournament may result in three possible outcomes: *greater than*, *lesser than*, and *equal to*. However, the authors also introduced a key difference that, as we will show later, it became problematic. When two solutions were evaluated to be equal, aiGA required a normalization step to transform the graph into another that only includes strict comparison relations; however that step was eliminated reducing the amount of comparisons stored by the partial ordering graph. That is, if two solutions are evaluated as *equal*, no edge between the solution vertex was introduced, which lead as we will show later to graphs not dense enough to provide an accurate enough reconstruction of the global order—as presented in section 2.1.

Results using this approach were worst and more unstable than the ones using the original steady-state scheme (Llorà et al., 2006). Next section presents introduce graph-theoretic analysis helps explaining why the incremental approach was condemned to fail whereas the original steady-state approach was able to reconstruct, accurately enough, the global order required to build the synthetic fitness.

3 Interaction Sizing in aiGAs

A key element to the success of the aiGA, as shown in Llorà et al. (2006), is its ability to reconstruct a global ordering out of the partial-order graph built using the user evaluations. However, there is little knowledge about what make a partial-order graph $G = \langle V, E \rangle$ a good candidate to generate an accurate enough global ordering, or ranking, of the solutions contained in V . Moreover, when facing the design of aiGA there are three variables that need to be taken into account: the length of the solutions (ℓ), the growth ratio of nodes ($\Delta(V)$), and the growth ratio of edges ($\Delta(E)$). Usually, given a certain problem to solve ℓ is fixed, but we then need to decide how often we need to generate a new node or show a new comparison to the user.

To illustrate the interaction-sizing dilemma, we will review two basic strategies. Two opposite strategies could be (1) to generate a new vertex (solution) every time we require a comparison from the user, and (2) only generate a new vertex (solution) when there is no comparison left between the one already present in G . From the user perspective, the first one will always provide new elements to compare alleviating the burden of repetitiveness problem, but would provide a poor global ranking. On the other hand, the second one will increase the repetitiveness perception and fatigue of the user, but completely connected graph provide a perfect estimate to reconstruct the global ranking of the solutions.

Also, these strategies will lead to graph with very different structural properties. Whereas the first one will form graph with a large number of unconnected components (each of them form by 2 vertexes), the second one will eventually lead to fully connected graphs—structurally similar to the proposed incremental method in Llorà et al. (2005). It is important to mention here, that there is a lower bound on the number of initial vertexes required¹. However, satisfying that lower bound does

¹As shown elsewhere (Cristianini and Shawe-Taylor, 2000; Shawe-Taylor and Cristianini, 2004; Llorà et al., 2005), the number of initial nodes is proportional to the number of training example required to properly train the ϵ -SVM regressor used to construct the synthetic fitness. A detailed explanation is beyond of this paper, been explained

still not guaranty that resulting graph G will be useful, since that is the results of the interactive policy selected.

For the above-mentioned reasons, we focused on identifying a graph-theoretic measure (Harary, 1971; Clark and Holton, 1991) that could shed some light on what is a good graph structure for the purposes of computing the global ranking. Since there seams to be a tradeoff between the number of nodes ($|V|$) and the number of edges ($|E|$), the graph *density* provides a proper unified measure across different graphs.

Definition 1 *The graph density for a directed graph is defined as:*

$$D(G) = \frac{|E|}{|V|(|V| - 1)} \quad (1)$$

Whereas the density for a non-directed graph is defined as:

$$D(G) = \frac{2|E|}{|V|(|V| - 1)} \quad (2)$$

Following the theoretical framework setup in Llorà et al. (2005), we also used a *perfect user*, as defined below, to conduct the analysis of the right graph structure required for the aiGA. This assumption will then be revisited in the next section.

Definition 2 *A perfect user Υ is a function that given 2 solutions u, v defined as:*

$$\Upsilon(u, v) = \begin{cases} \langle v, u \rangle & \text{if } f(u) < f(v) \\ \langle u, v \rangle & \text{if } f(u) > f(v) \\ \lambda & \text{if } f(u) = f(v) \end{cases} \quad (3)$$

where f is the ideal target function, and λ stands for the null edge.

The rest of this section assumes that f is the OneMax function (Goldberg, 1989) following the original aiGA facet-wise analysis (Llorà et al., 2005). Also, assuming the usage of a *perfect user*, the density definition presented in Definition 1 for a partial-order graph can be reformulated as:

Definition 3 *The graph density for a partial-order graph given a perfect user Υ is defined as:*

$$D(G_{\Upsilon}) = \frac{2|E|}{|V|(|V| - 1)} \quad (4)$$

When using a *perfect user* Υ , Definition 3 presents the same density form of 2. This is the result of the Υ behavior. Given 2 nodes u and v the will only return one of the two possible edges $\{\langle u, v \rangle, \langle v, u \rangle\}$.

The density of the graph generated by the hierarchical tournament method used by the original aiGA—see figure 1—requires computing the number of vertexes and edges after each round i of tournaments. The number of nodes nn in a partial-order graph at a given tournament round i of height h can be computed as:

$$nn(i, h) = 2^h + i \cdot (2^{h-1} - 1) \quad (5)$$

Whereas the number of edges ne at a given tournament round i of height h can be computed as:

$$ne(i, h) = (i + 1)(2^h - 1) \quad (6)$$

elsewhere (Llorà et al., 2005).

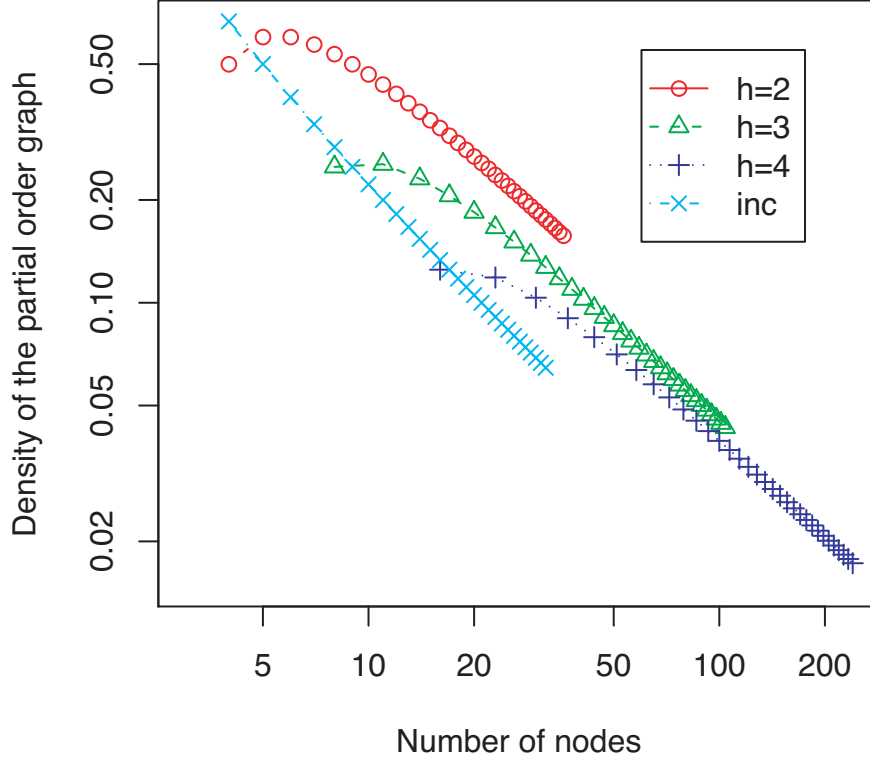


Figure 2: Density of the partial-order graph create by a hierarchical tournament scheme given the height h and the tournament round i . The figure plots the densities for $h = \{2, 3, 4\}$ and $i \in [0, 32]$. Figure highlights the quick drop on density on the incremental approach.

Hence, the density of the partial-order graph at a given tournament round i of height h can be expressed as:

$$D(G^{i,h}) = \frac{2 \cdot ne(i, h)}{nn(i, h)(nn(i, h) - 1)} \quad (7)$$

Figure 2 presents the density across different heights and iterations for the hierarchical tournament in the original aiGA. Using the same density measure, we revised the incremental policy proposed in Llorà et al. (2006). Figure shows a clear change in slope and drop in density, providing inaccurate global rankings. It also does not allow to scale based on the problem size ℓ , violating the requirements of the ϵ -SVM regressor, which may lead to inaccurate regressions. Using the density measure we were able to explain why the incremental tournament approach proposed in Llorà et al. (2006) did not work. However, this does not answer why the hierarchical tournament scheme is able to ensemble partial-order graphs that provide accurate global rankings. In order to be able to answer such a question, we prepared a controlled experiment in order to characterize the boundary separating accurate and inaccurate partial-order graphs.

We can label a partial-order graph as an accurate one if and only if, after optimizing the synthetic fitness produced, the best solution obtained has correctly fixed at least $m - 1$ building blocks (Goldberg, 2002). In the particular case of solving OneMax, the best solution generated has at least all but one bit set to 1, otherwise, we will label such a partial-order graph as inaccurate. In order to determine where that boundary existed, we proceed as follows. Given a certain problem length ℓ , we generate 75 random graph for each of the possible number of nodes and edges. That is, for $\ell = 2$ we would generate random graph containing $\{1, 2, 3, 4\}$ nodes, and for each node number

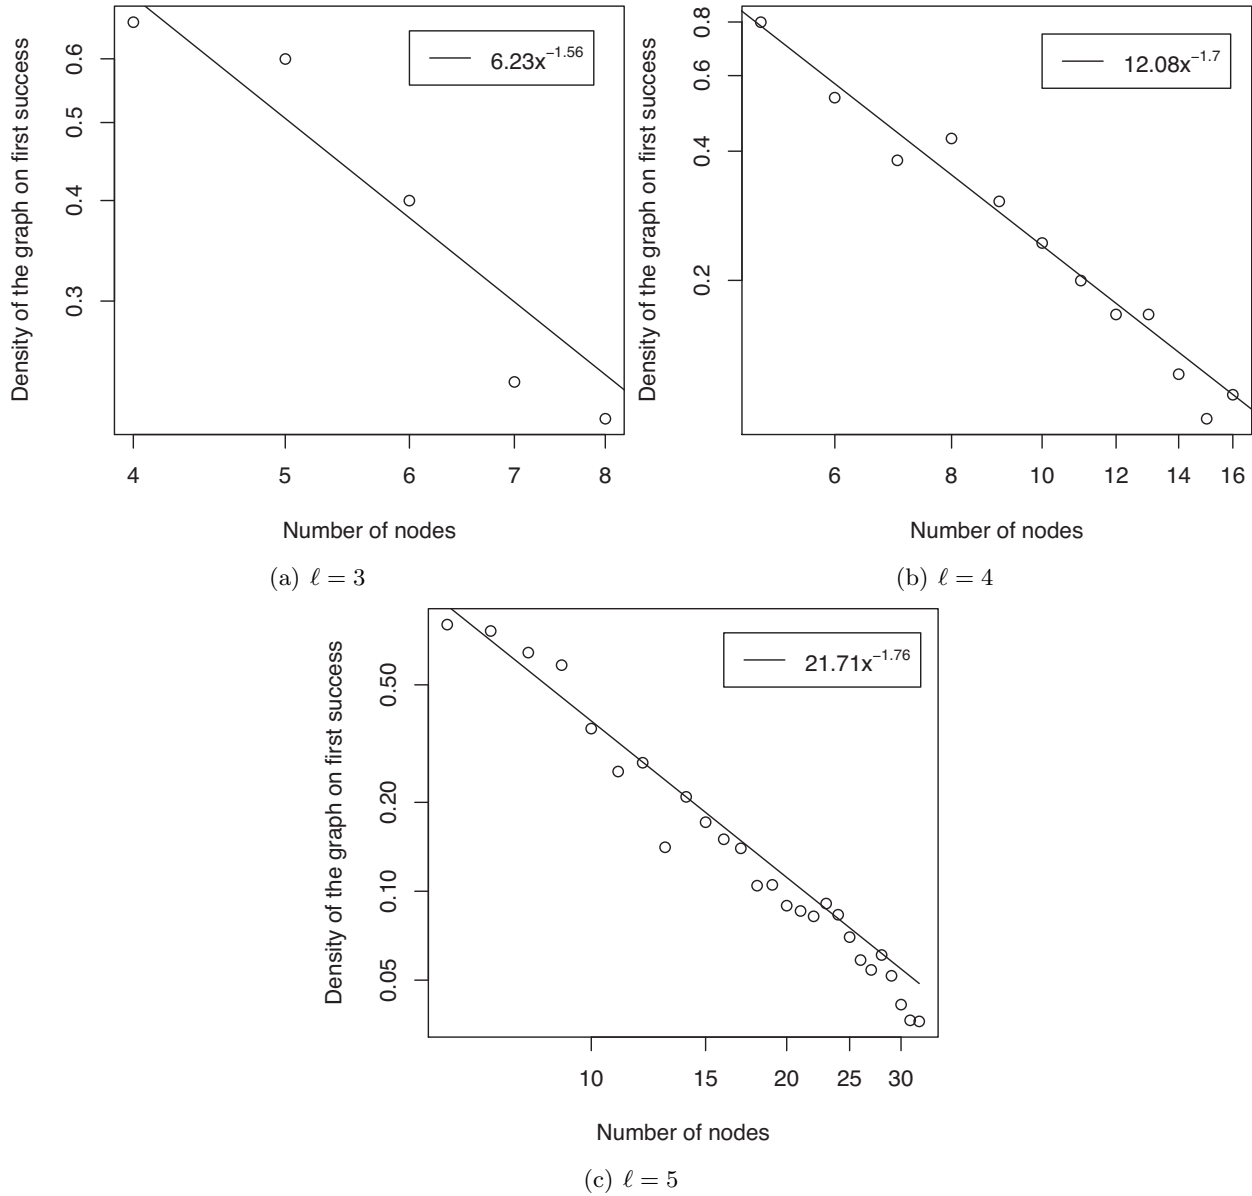


Figure 3: Figures presents the density required to hit the optima using a randomly generated graph with the given number of nodes. Each point represents the average density of 75 successful independent trials—all random graphs induced a solution after the synthetic fitness optimization which had correctly fixed at least $m - 1$ building blocks.

we would generate random graphs for all the possible number of edges $1 \rightarrow \{\}, 2 \rightarrow \{1\}, 3 \rightarrow \{1, 2, 3\}, 4 \rightarrow \{1, 2, 3, 4, 5, 6\}$. Hence, for $\ell = 2$ we generated $75 \cdot (1 + 1 + 3 + 6) = 825$ random graphs. For each of them we computed their density and the probability of success (proportion of accurate partial-order graphs). Figure 3 shows given a number of nodes what was the required density to achieve for the first time an accurate partial-order graph—all the 75 graphs for that density were accurate. After achieving an accurate graph for the first time given a certain number of nodes, increasing the density just keeps generating accurate graphs.

Figure 3 describes a lower bound on the required density of the partial-order graph required to ensemble accurate partial-order graphs based on the user evaluations. Figure 3 and 2 needs to be carefully compared. There is only one case that is directly comparable, the case where $\ell = 4$ —Figure 3(b)—and we use a hierarchical tournament of height $h = 2$ —worst case scenario. The densities of the hierarchical tournament scheme of the aiGA (Llorà et al., 2005) and the boundary obtained align nicely, following the same power law— $D(G_{HTS}) \approx D(G_{rnd})$. The tail between $\ell = 4$ and $h = 2$ eventually reverse the balance, but on an interactive GA session that would happen far after the user has finished the experiment exhausted. If we increase the height by 1, then the produced partial-order graph densities are higher than required but will guaranty the creation of an accurate partial-order graph. Hence, the hierarchical tournament (Llorà et al., 2005) unexpectedly turned out to be closer to the density boundary required to produce accurate partial-order graphs, and that is also explains why the incremental approach (Llorà et al., 2006) did not work—it did not satisfy the lower density boundary.

4 Parallel Evaluation Ensemble in aiGAs

Another challenge for any interactive genetic algorithm is how to integrate the results obtained from different runs. If you are trying to solve a problem by means of an iGA, you may end having multiple run results that you need to combine. Moreover, those results may also be coming from different users and, thus, it is highly likely that results may differ. Traditional approaches try to combine (for instance averaging) the solutions obtain from the different interactive sessions. Formiga and Alías (2007) have mostly focused on combining the end products of the interactive evolution, because that was what was left after the execution. However, aiGAs beside providing a handful of good candidate solutions at the end of the run, they also provides the partial-order graph. Such information can be of great value. As shown in Alías et al. (2006), partial-order graphs can help eliminate user contradictions by guided reevaluation of cycles—a sign of user contradictory evaluation criteria.

Thus, our approach to generate a final answer to the problem to be solved after n interactive session will involve assembling the n partial-order graphs to create a final synthetic fitness function. The promising candidates resulting from the synthetic fitness function will form the answers to the overall problem. The goal is to benefit from all the embedded knowledge in those partial-order graphs looking for a robust methodology—combining final solutions has been shown to produce average solution of poor quality (Formiga and Alías, 2007). Figure 4 illustrates how such an ensemble may be constructed.

The process of building the ensemble \hat{G} from a set of partial-order graphs $\{G_0, G_1, \dots, G_n\}$ has to consider different sources of possible contradictions. In another words, what could be the different causes of having edges such as $\langle u, v \rangle$ and $\langle v, u \rangle$ competing to enter the final ensemble \hat{G} . A set of possible, but not all, causes is listed below.

- **User mistakes:** During the interactive process the user just got distracted and provided the wrong tournament evaluation adding a contradictory evaluation.




Figure 4: Creating a new partial-order graph by combining the graphs obtained in interactive sessions using aiGA. Conflicts may arise because of contradictory evaluations—see edge between b and c . Those conflicts must undergo a conflict resolution policy.

- **Perceptually indistinguishable:** The interface of the interactive process is not refined enough. Hence, given two solutions u and v , the user is not able to perceptually distinguish the difference between them. This situation increases the risk of providing contradictory evaluations.
- **Contradictory targets:** When dealing with aesthetic evaluations, the objective target that each user may have in mind can clearly differ. A simple analogy would be a multimodal problem where two radically different solutions provide the same fitness. In such situations, blindly enabling the partial-order graph will lead to mediocre average solutions.
- **Criteria shifts:** The interactive nature of the process can also shift the user criteria. Such criteria shifts tend to appear in long lasting interactive sessions (Alm and Llorà, 2006). Users can change their desired target based on the interactive process itself. That is, user may find a new point of interest that suddenly became much more interesting than the original one, producing the criteria shift.

Addressing all these possible sources of user evaluation contradictions is beyond the reach of this paper, and will require further research. However, the rest of this section will focus only on the first source of possible evaluation contradictions: *user mistakes*. As mentioned earlier, during the interactive aiGA process the user just got distracted and provided the wrong tournament evaluation adding a contradictory evaluation. Such kind of mistakes can be easily modeled as a binomial distribution. Let p_{er} be the probability of edge reversal due to a user mistake in the evaluation process. Thus, for each pair of adjacent solutions (vertex) u and v , we can model the probability of edge reversal ($e = \langle u, v \rangle$ and $e = \langle v, u \rangle$ the reversed edge) as a binomial distribution $B_e(n, p)$, where n is the number of graph G_i that contain e or e and p the probability of user mistake—usually p is small thanks to the users efforts ($p < 0.05$). Thus, the expected number of reversed edges can be written as np . Since p is a small number, in order to maximize the probability of choosing between e and e the right edge, a simple strategy would be to pick the one with larger support—usually know as majority rule (Mitchell, 1997; Witten and Frank, 2000).

5 Conclusion

This paper has shown a first analytic study based on graph-theoretic measures to separate accurate partial-ordering graphs (the ones that can provide an accurate ranking that will lead to accurate solutions) from the inaccurate ones. As a result, we have identified a lower bound based on the density of the partial-order graph. If the partial-order graph density is below that boundary, the algorithm will not be able to provide high quality solutions. Previous failed efforts conducted by researchers to replace the original hierarchical steady-state tournament evaluation proposed by aiGAs with incremental versions, can now be explained. In order to be able to create an accurate global ranking using a partial-order graph, such a graph requires a density above the lower bound identified.

In this paper we have also addressed the problem of addressing the creation of partial-order graph ensembles based on graph provided from previous aiGA sessions. We have focused on how we can model user mistakes and provide a strategy to build accurate partial-order ensembles. We showed that such mistakes can be probabilistically modeled as binomial distributions and, thus, simple contradiction resolution policies can be put in place to maximize the quality of the resulting ensemble based on the majority rule.

6 Acknowledgments

This work was sponsored by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF, under grant FA9550-06-1-0096, the National Science Foundation under ITR grant DMR-03-25939 at Materials Computation Center and under grant ISS-02-09199 at the National Center for Supercomputing Applications, UIUC. The U.S. Government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation thereon.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research, the National Science Foundation, or the U.S. Government.

The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research, the National Science Foundation, or the U.S. Government.

References

- Alías, F., Llorà, X., Formiga, L., Sastry, K., and D.E., G. (2006). Efficient interactive weight tuning for TTS synthesis: reducing user fatigue by improving user consistency. In *Proceedings of ICASSP*, volume I, pages 865–868, Toulouse, France.
- Alm, C. O. and Llorà, X. (2006). Evolving emotional prosody. In *Proceedings of the Ninth International Conference on Spoken Language Processing (INTERSPEECH 2006)*, page paper 1741.
- Caldwell, C. and Johnston, V. S. (1991). Tracking a criminal suspect through face-space with a genetic algorithm. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 416–421. Morgan Kaufmann.
- Clark, J. and Holton, D. A. (1991). *A First Look at Graph Theory*. World Scientific.

- Coello-Coello, C. A. (December, 1998). An updated survey of GA-Based Multiobjective Optimization Techniques. Technical report lania-rd-09-08, Laboratorio Nacional de Informática Avanzada (LANIA), Xalapa, Veracruz, México.
- Cristianini, N. and Shawe-Taylor, J. (2000). *An Introduction to Support Vector Machines*. Cambridge Press.
- Dawkins, R. (1986). *The blind watchmaker*. New York: W. W. Norton.
- Deb, K., Agrawal, S., Pratab, A., and Meyarivan, T. (2000). A Fast Elitist Non-Dominated Sorting Genetic Algorithm for Multi-Objective Optimization: NSGA-II. KanGAL report 200001, Indian Institute of Technology.
- Formiga, L. and Alías, F. (2007). Extracting user Preferences by GTM for aiGA Weight Tuning in Unit Selection Text-to-Speech Synthesis. *Lecture Notes in Computer Science*, (4507):654–651.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley.
- Goldberg, D. E. (2002). *The design of innovation: Lessons from and for competent genetic algorithms*. Kluwe Academic Publisher.
- Goldberg, D. E., Korb, B., and Deb, K. (1989). Messy genetic algorithms: Motivation, analysis, and first results. *Complex Systems*, 3(5):493–530.
- Harary, F. (1971). *Graph Theory*. Addison-Wesley Publishing Company.
- Harik, G., Lobo, F., and Goldberg, D. E. (1998). The compact genetic algorithm. *Proceedings of the IEEE International Conference on Evolutionary Computation*, pages 523–528. (Also IlliGAL Report No. 97006).
- Llorà, X., Goldberg, D. E., Ohsawa, Y., Matsumura, N., Washida, Y., Tamura, H., Yoshikawa, M., Welge, M., Auvil, L., Sears-Smith, D., Ohnishi, K., and Chao, C.-J. (2006). Innovation and Creativity Support via Chance Discovery, Genetic Algorithms, and Data Mining. *Journal of New Mathematics and Natural Computation*, 2(1):85–100.
- Llorà, X., Sastry, K., Alías, F., Goldberg, D. E., and Welge, M. (2006). Analyzing active interactive genetic algorithms using visual analytics. In *GECCO '06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 1417–1418, New York, NY, USA. ACM.
- Llorà, X., Sastry, K., Goldberg, D. E., Gupta, A., and Lakshmi, L. (2005). Combating user fatigue in iGAs: Partial ordering, support vector machines, and synthetic fitness. In *Proceedings of Genetic and Evolutionary Computation Conference 2005 (GECCO'05)*, pages 1363–1370, Washington DC, USA. ACM.
- Mitchell, T. M. (1997). *Machine Learning*. McGraw Hill.
- Sastry, K. (2001). Evaluation-relaxation schemes for genetic and evolutionary algorithms. Master's thesis, University of Illinois at Urbana-Champaign, General Engineering Department, Urbana, IL. (Also IlliGAL Report No. 2002004).
- Shawe-Taylor, J. and Cristianini, N. (2004). *Kernel Methods for Pattern Analysis*. Cambridge Press.

- Takagi, H. (2001). Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. In *Proceedings of the IEEE*, volume 89, pages 1275–1296.
- Witten, I. H. and Frank, E. (2000). *Data Mining: practical machine learning tools and techniques with java implementations*. Morgan Kaufmann.