
The innovation pump: supporting creative processes in collaborative engineering

Xavier Llorà*

Data-Intensive Technologies and Applications,
National Center for Supercomputing Applications,
University of Illinois at Urbana-Champaign,
1205 W. Clark Street, Urbana, IL 61801, USA
E-mail: xllora@illinois.edu
*Corresponding author

David E. Goldberg

Illinois Genetic Algorithms Laboratory,
University of Illinois at Urbana-Champaign,
104 S. Mathews Avenue,
Urbana, IL 61801, USA
E-mail: deg@uiuc.edu

Abstract: The pervasive expansion of computers and internet has change the way people collaborate. Cybercollaboratories for collaborative engineering in form of web boards, blogs, e-mails, and instant messaging have become de facto mainstream communication channels. This paper reviews the new framework set after these technologies and presents how collaborative creativity and innovation can be modelled and supported using computational models. The paper continues presenting an innovation-support model based on the usage of genetic algorithms as computational metaphors of human innovation. The paper also discusses the results achieved using the proposed technologies in real-world collaborative creative processes.

Keywords: creativity and innovation; computer-supported innovation; collaborative creative processes; interactive genetic algorithms; models of human innovation.

Reference to this paper should be made as follows: Llorà, X. and Goldberg, D.E. (2009) 'The innovation pump: supporting creative processes in collaborative engineering', *Int. J. Collaborative Engineering*, Vol. 1, Nos. 1/2, pp.75–97.

Biographical notes: Xavier Llorà's work on Genetics-Based Machine Learning (GBML) has placed him among the leaders of the renaissance of learning classifier systems. He has recently edited two books reviewing the advances of the field from 2003 to 2009, and served as GBML track chair in the ACM SIGEVO GECCO conference. Since his arrival to University of Illinois at Urbana-Champaign in 2003, he has headed the DISCUS project to support collaborative human-innovation and creativity with several national and Japanese universities and consortiums as partners. In 2005, he was named research assistant professor and became research staff at the National Center for Supercomputing Applications.

David E. Goldberg's first book arguably did as much any single reference to spur worldwide growth in al genetic and evolutionary. Goldberg started his academic career at the University of Alabama, Tuscaloosa. He was awarded a 1985 NSF Presidential Young Investigator Award, and moved to the University of Illinois at Urbana-Champaign in 1990, where he established. The Illinois Genetic Algorithms Laboratory (IlligAL). In 2003 he was named Jerry S. Dobrovolny Distinguished Professor in Entrepreneurial Engineering. He is founder, fellow, and past chairman of the International Society for Genetic and Evolutionary Computation, which recently merged with the ACM to become SIGEVO.

1 Introduction

It is hard to imagine nowadays any engineering activity that does not heavily rely on computer-mediated communications. The pervasive penetration of computers and internet into almost any facet of human activities has change the way people collaborate. Terms such as cybercollaboratories and cyberinfrastructure are getting traction in day-to-day work. Web boards, blogs, e-mails, and instant messaging have become, *de facto* mainstream communication channels. People scattered across the globe collaborate thanks to such technologies to carry out their daily work. The massive adoption of such technologies for collaborative communication has lead to a new scenario: the explosion of the amount of information available. E-mails or instant messaging logs, for instance, also provide fossilised records of group communications during collaborative endeavours.

Creativity-driven processes – a key player in collaborative engineering-have also taken advantage of such new communication media. This paper briefly reviews the new framework set after these technologies. Any collaborative creative process involves, at least, dealing with large volumes of archived information, collaboration data provided by computer-mediated communication channels, and – in the collaborative engineering case – simulations or descriptions of the solution for the given problem under discussion. This cloud of digital data hides relevant knowledge about creative processes that, if properly used, unveils new hints and guidance to assist such a creative process setting a new parading for collaborative engineering. It also highlights the urgent need for creativity and innovation support technologies that assists the participants to navigate and reflect in a computer-mediated collaborative endeavour.

The work presented in this paper is the result of the research conducted in the *Distributed Innovation and Scalable Collaboration in Uncertain Settings* (DISCUS) project¹ (Goldberg et al., 2003). Using Genetic Algorithms (GAs) as a metaphor of human-innovation, the work presented in this paper describes how learned computational models of user preferences provide a natural aid to fuse the human-computer creativity endeavour. Active interactive GAs (Llorà et al., 2005) can help guide the participants of a problem-solving oriented task to better navigate and reflect across the creative endeavour. In another words, active interactive GAs act as an *innovation pump* to aid the participant providing him with *educated guesses* about the problem solution as a form of *knowledgeable recombination*. The paper concludes relating some of the efforts conducted under the DISCUS project to apply such an innovation pump to collaborative engineering real-world problems.

The rest of the paper is structured as follows. Section 2 reviews the motivations and need for creating innovation and creativity support technologies for collaborative endeavours, and how collaborative engineering may directly benefit of the created technologies. Then, Section 3 presents the decomposition of the main components involve in the creation of an *innovation pump*. Once the key elements are introduced, Section 4 describes the creation of an *innovation pump*, later used on real-world engineering problems, as described in Section 5. Finally, Section 6 presents some conclusions about the work presented in this paper.

2 Innovation and creativity for collaborative engineering

2.1 A whole new game at your finger tips

Modern times challenge organisations and their leaders to adapt quickly and well to complex, fast-moving circumstances under trying conditions. Data sources are numerous, distributed, and contradictory. Challenges are difficult to detect and diagnose, widely dispersed, and constantly changing. Sources of knowledge and expertise are distributed, of varying quality, and difficult to integrate. Moreover, the tools of the trade are increasing in technological sophistication, computational intensity, and require specialised hardware, software, and care and feeding.

As a result, many have sought to build combinations of information technology under the rubric of *Knowledge Management* (KM) to support collaboration and the integration of multiple data sources. Loosely defined, KM integrates IT and people to improve organisational learning, collaboration, and improvement. KM initiatives may be as simple as building databases of organisational competence or they may involve integration of ones customers directly to production, marketing, engineering, and product planning staff. But first-generation KM looks like a simple extension of *batch processed* management information systems of the 1970s and 1980s, and even when continuous improvement is involved it consists of a slow-moving single loop (Goldberg et al., 2003).

Modern challenges demand a more integrated, interactive, and evolvable process to make the most of human and computational inputs to better answer the challenges post by complex environments. Moreover, collaborative endeavours for creative tasks could greatly benefit from the massively adopted mainstream computer-mediated communication channels. For instance, real-time analysis of web forums and interactive manipulation of computer models and simulations are only the tip of the iceberg. Collaborative engineering can take a whole new dimension due to the tight integration of simulation, collaboration tools, and access to vast archives with a few clicks.

2.2 The online medium challenge

A primary difficulty of using online communication for serious innovation and creativity is its superficiality. Compared to in-person communication with its relatively high emotional affect or written communication with its thoughtfulness born of permanence, online communication can be relatively superficial, making it difficult to reflect on deeper issues. For this reason, any attempt to provide a systematic approach to innovation based on an online media requires to provide effective means for human reflection and reasoning. Meaningful visual representations of discourse and

interaction – human-human and human-computer – is a easy way to provide the users with the basic tools that will help them focus on the discourse goals. Moreover, thanks to the computer-mediated interaction, online information can be stored and mined to identify main interaction building blocks and patterns between the participants – as we later present.

2.3 *Integration in a cloud of components*

The work presented in this paper is the results of the basic research in the creation of a novel IT system under the DISCUS project. The goal of DISCUS is to create a distributed and scalable environment for the integration of both human- and computer-generated knowledge in uncertain settings through effective human-human and Human-Machine Collaboration (HMC). Normal human collaboration and innovation involves Human-Human Collaboration (HHC) through cross-fertilising exchange between pairs of larger teams of individuals. The use of the tools of modern artificial and computational intelligence by individuals is a kind of HMC. To achieve effective and scalable collaboration in difficult problem domains, this project envisions both types – HHC and HMC – as being important.

To achieve these goals, the DISCUS system proposes the integration of a number of key elements:

- Interactive Genetic Algorithms (iGAs)
- Human-Based Genetic Algorithms (HBGAs)
- Scalable Genetic Algorithms (machine based)
- Flexible data and text mining (D2K/T2K)
- Chance discovery using Key Graphs.

The remainder of this subsection briefly discusses each of these elements. Key targets of basic research opportunities will be discussed in more detail later.

Interactive Genetic Algorithms. GAs are search procedures based on the mechanics of natural selection (Goldberg, 1989). They combine

- general and independent evaluation of solution quality or merit
- the coding of the set of possible solutions as a set of alternative chromosomes or genotypes
- the selection of better solutions according to merit
- genetic-like variation mechanisms such as crossover and mutation to promote the rapid generation of new, possible better, solutions to a user's problem.

GAs have increasingly been used across the spectrum of human endeavour (Goldberg, 1994). iGAs are those that replace the computer computation of the relative fitness of solutions, or their objective function values, with the judgement of a human evaluation.

Dawkin's Blind Watchmaker program (1986) and the Faceprints system developed at New Mexico State University (Caldwell and Johnston, 1991) are two early examples of

iGAs. For example, in Face prints, 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 sub codes represent different facial features (nose type, mouth type, hair type, etc.). Each full chromosome maps to a face and the population of chromosomes 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. Takagi (2001) presented detailed review of the progress of interactive GAs and interactive Evolutionary Computation (iEC).

Human-based Genetic Algorithms. Whereas iGAs replace the computerised evaluation of an objective function with the human evaluation of what might be termed a subjective function, human-based GAs (Kosorukoff and Goldberg, 2002) go further and permit evaluation, selection, and variation to be performed by a human. A simple example of such a model brings together users with problems and invites solutions to those problems in the form of text responses. Users are encouraged to select the solutions of others or create their own. Many 'new' solutions are simple combinations or elaborations of previous solutions on the list.²

Scalable Genetic Algorithms. GAs are a core technology of this project, but recent work (Goldberg, 2002) has shown that first-generation GAs – those in common use – do not scale well on difficult problems. This does not imply that simple GAs is not useful; however, it does suggest that more scalable GAs should be designed and used. The past decade has seen great strides in the design of competent GAs – GAs that solve a large, important class of hard problems quickly, reliably, and accurately. Procedures such as the fast messy genetic algorithm (Goldberg et al., 1993) and the hierarchical Bayesian optimisation algorithm (Pelikan and Goldberg, 2001) are able to solve what Simon (Simon, 1969) called nearly decomposable problems in times that grow no more quickly than a quadratic function of the number of decision variables in a problem.

Flexible Data Mining. DISCUS is the product of ongoing collaboration between the Automated Learning Group (ALG) of the National Center for Supercomputing Applications (NCSA) and the Illinois Genetic Algorithms Laboratory (IlligAL). ALG has developed a tool called D2K – Data to Knowledge³ (Welge et al., 2003) – a visual programming environment for building data mining applications. D2K is a highly flexible environment that allows developers and analysts to build data and/or text mining applications while incorporating human-generated knowledge, computer-generated knowledge and chance discovery into the process. The knowledge extracted through the D2K/T2K analysis will provide the decision maker with a set of understandable patterns. These understandable patterns can be used to

- make predictions and classification about new events
- form the basis for the discovery of a new hypothesis
- detect and monitor outliers in databases
- summarise the contents of large data stores.

All of these possible outcomes support the DISCUS process.

Chance Discovery. Statistical techniques for machine learning and data-mining implicitly rely on sufficiently large data samples to give the researcher confidence that the knowledge mined is statistically significant; however, human mining of data often relies on few samples and a heightened sensitivity to and interpretation of unusual occurrences or chance events. The literature of scientific discovery is filled with poignant examples (the discovery of penicillin and Post It notes are two examples that come to mind), and with this in mind, there is an intellectual movement afoot to understand the role of chance discovery (Ohsawa, 2002).

Among the computational tools being developed to discover and exploit chance occurrences is the so-called Key Graph. Key Graphs are graphical representations of text and other delimited data that group

- statistically prominent words and their relationships to others
- statistically less frequent words and their relationships to clusters of the prominent words.

While statistically prominent words are often the focus of traditional data mining, humans familiar with the situations represented by the data often find statistically prominent clusters to be ‘obvious’. A critical insight of chance discovery is to dig beneath the prominent clusters and mine less prominent, but possibly meaningful terms.

Another key chance discovery technique is the Influence Diffusion Model (IDM) (Matsumura, 2003; Matsumura et al., 2005). Tracking key terms in a creative discussion may help the participants on an online discussion to stay focus on the goal of the discussion. However, several other important facets can be also analysed. IDM measures the influence of comments, participants, and terms by the degree of text-based relevance of the comments when interactions among participants are done by exchanging comments. The threaded comments, called *comment-chain*, show the flow of influence across the discussion. Moreover, IDM also provides a measure of the influence of terms, participant, and the social network of influence between them.

2.4 *The innovation intuition as a metaphor of human innovation*

GAs are a core technology for the creation of an innovation support infrastructure. Starting in 1983, Goldberg (1983, 2002) developed the so called *fundamental intuition of genetic algorithms*, or the *innovation intuition*. Specifically, the innovation intuition of GAs is about the work together of:

- selection and mutation
- selection and recombination.

Moreover, the innovation intuition of GAs provide a facet-wise modelling of human innovation. This approach models two orthogonal *facets of human innovation*.

Selection + mutation = Continual improvement. Selection and mutation working together are a form of hill-climbing mechanism. Mutation suggests variants in the neighbourhood of the current solutions; selection acts as the decision process, which accepts improving changes with a high probability. This simple model describes one of the facets of

human innovation, the so called *continual improvement* in total quality management literature, or as Japanese call it, *kaisen*.

Selection + crossover = Innovation. Another facet of human innovation is the so-called *cross-fertilising innovation*. People usually grasp a set of good solution features in one context, and a notion in another context and juxtaposing them, thereby speculating that the combination might be better than either notion taken individually. Taking together selection and crossover, GAs are a computation model of *cross-fertilising innovation*.

The *innovation intuition* as a metaphor of human innovation is key to the *innovation technology* revolution. As early mentioned, humans are to become the main measure of such a technology. Pervasive GA-guided interaction between human and computers opens a new research path to creativity- and innovation-support. Two well-known models of such support are interactive GAs, and human-based GAs. Interactive GAs iGAs replaces the computer computation of the relative fitness of solutions and the selection process by the judgement of a human evaluation. More detailed information about the progress of interactive GAs and iEC are presented in a review by Whereas iGAs replace the evaluation and selection by the human judgement, human-based Gas (Kosorukoff and Goldberg, 2002) HBGAs move one step further and permit evaluation, selection, and variation to be performed by a human. For such reasons, the previous facets of GAs may be regarded as a first order model of human innovation – as we will later discuss in this paper.

3 A bare-bones approach

This section describes how the approach proposed under the DISCUS project can be applied to the collaborative creation of marketing scenarios using focus groups. Even a simple model (Llorà et al., 2004a) relies on mechanisms to support innovation and creativity via – among other components – iGAs. This section also introduces the key challenges when using interactive GAs as the engine of an *innovation pump*.

3.1 The collaborative marketing scenario

DISCUS has been mainly applied to the collaborative creation of marketing scenarios (Llorà et al., 2004a, 2006a) using focus groups. The participants in creative processes have access to enhanced collaboration tools including message boards, instant messaging capabilities, and chat rooms using a simple browser – detailed descriptions may be found somewhere else (Llorà et al., 2006b). DISCUS organises the participants in discussion groups. A discussion owner is in charge of managing the group and deciding the finalisation of the discussion. DISCUS allows parallel and sequential discussions. Besides the information exchange inside a group discussion, information may be shared among group discussions, regardless of their sequential or parallel interaction, using *solution centres*. The outcome of group discussion is archive in a solution centre.

Using the proper access and time policies, they become the proper channel to pass information around the different discussion groups.

These *solution centres* are a key element of DISCUS. They contain all the proposed and analysed solutions – or simulations in some collaborative engineering setups – for the given problem under discussion. A key element to aid user's creativity is the ability to retrieve and generate promising solutions. The rest of this section, after a brief overview of the DISCUS setup for the collaborative creation of marketing scenarios, focus on introducing the main elements involved in the usage of iGAs, the key component of the *innovation pump* proposed in this paper.

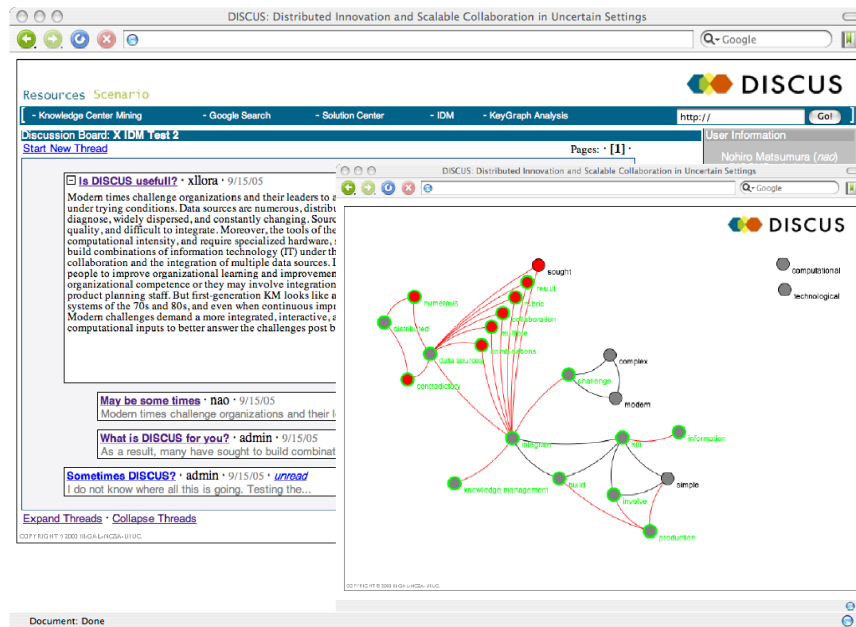
3.2 *All components in action*

After three initial experiments to validate the DISCUS concept (Llorà et al., 2004a, 2004b, 2004c) the platform was ready for a real-world experiment in the marketing research area. March 2005 became the achievement of a big milestone for the DISCUS project (Llorà et al., 2006b). DISCUS researchers help HakuHodo Inc. to discover emerging markets and sketch new products that will satisfy coming needs. The goal of the experiment was to identify future scenarios for cell phone usage and the features that will make them popular among consumers.

The March experiment used DISCUS to streamline the data analysis, focus group brainstorming, and scenario creation and evaluation. The data gathered by the marketing firm – questionnaires about cell phone usage completed by people in New York, Los Angeles, and Chicago – was mined using the D2K components of DISCUS in order to determine what characteristics the team should look for in focus group participants. Long-time tested marketing diffusion models provided by the marketing researchers guided this data mining process. Then, the DISCUS team conducted a similar survey on the Urbana-Champaign campus of the University of Illinois. Surveys provided by over a hundred respondents were analysed; 50 respondents were recruited according to the desired characteristics to participate in the focus groups. The goal was to have a combination of technology innovators, who are power users of gizmos; early adopters; and late adopters. Different profiles of focus groups were created based on these criteria.

Focus groups were conducted using DISCUS collaborative software on desktop computers to answer questions, participate on free form discussion, and provide feedback. While participants in the focus groups discussed, communications were analysed and plausible future scenarios designed by marketing and DISCUS researchers. Scenario identification by means of Key Graphs helped the researcher to discover promising future scenarios in cell phone usage. Moreover, influence diffusion among the social network of the participant in the focus group was also available during the discussion. Such influence diffusion analysis helps researchers to identify communication roles in focus groups, and easily redirect the process by influencing the communication leaders. The identified roles sometime clearly differ from the original intents due to the nature of the participants. Figure 1 shows one focus group where a new and unexpected moderator emerge during the discussion.

Figure 1 An ongoing DISCUS group discussion. The graph visualisation presents the main topics discussed and their relations. It acts as a visual shared dashboard for group reflection (see online version for colours)



3.3 Supporting scenario creation via innovation and creativity support technology

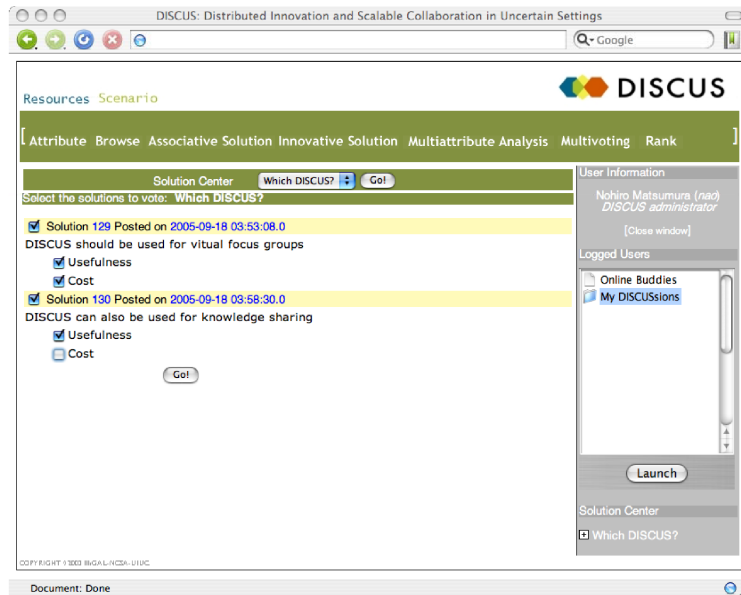
The creation of new marketing scenarios (Llorà et al., 2004a, 2006) requires the users to evaluate and select relevant solutions. Supporting user's innovation would require providing new solution as an innovative recombination of the main components of the existing ones. For this purpose, interactive GAs is more than a metaphor. As earlier introduced, they combine

- general and independent evaluation of solution quality or merit provide by the subjective evaluation of the user
- the coding of the set of possible solutions as a set of alternative chromosomes merit
- genetic-like variation mechanisms such as crossover and mutation to promote the rapid generation of new, possible better, solutions to a user's problem.

In the generation of new marketing scenarios, DISCUS (Llorà et al., 2006b) provided a basic capability of interaction for allowing the participants to engage such interactive process. Participants, under the form of voting, selected the most appealing scenarios that better illustrate feasible solutions to the marketing discussion – see Figure 2. Thus, a simple bare-bones interactive genetic algorithm was used. Even in this primitive form where marketing experts did the recombination of ideas, such an aid represented a big boost in supporting the creativity and innovation of the focus group. For these reason, the next logical step was to replace the human intervention for creating an automated pump of innovation. For instance, in engineering environments computer simulations

guide the solution search. Each solution is usually described across a wide variety of decision variables. The effective usage of interactive GAs provided a natural fusion of human and computer innovation capabilities in the creative process.

Figure 2 An ongoing DISCUS group discussion. Participants, under the form of voting, select the most appealing scenarios that better illustrates feasible solutions to the marketing discussions (see online version for colours)



3.4 *The challenge of an innovation pump*

The uses of interactive GAs 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 optimisation. Takagi (2001) presented a review of research efforts related to the iGAs challenges. These research areas included:

- discrete fitness value input method
- prediction of fitness values
- interface for dynamic tasks
- acceleration of iGAs convergence
- combination of evolutionary and non-evolutionary computation
- active intervention
- theoretical research.

These areas may be reorganised in five main elements that iGAs need to address on their road to effective solution:

- *Clear goal definition.* A precise description of the goal is key element to help the user engage a successful innovation process. A clear definition helps evolve high-quality solutions. Moreover, if such definition is maintained along the run, the user's task gets greatly simplified.
- *Impact of problem visualisation.* The solutions presented to the user need to be understandable and comparable. If the visualisation is too complex, the user will be overwhelmed with details. If there is no simple way to qualitatively compare solutions, the user may not be able to make a proper decision. If such qualitative comparison is not easy, the quality of the user revaluations will decrease and greatly penalise the performance of the iGAs.
- *Lack of real fitness.* iGAs lack a quantitative fitness function analogous to the one used in traditional GAs. The qualitative nature of the evaluation process usually leads to scenarios where the user is asked to provide solution rankings or relative evaluations among a selected subset of solutions.
- *Fatigue and frustration.* User fatigue is a critical element to produce high-quality solutions. Long times until convergence, lead to tedious and demanding attention periods on the user side. Fatigue becomes the main reason of an early stop of the iGAs process and, hence, leads to low-quality solutions. Moreover, fatigue can easily lead to user frustration if appreciable results are not achieved.
- *Persistence of user criteria.* The user can change his evaluation criteria along an iGAs leading to a noisy evaluation scenario. The user criteria may drift along the run, leading to a dynamic optimisation scenario. Methodologies for helping the user to maintain the persistence of his evaluation criteria along the iGAs run are a key element.

The rest of the work presented in this paper focuses on the lack of a real fitness function and how we can take advantage of the relative evaluations provided by the user to reduce the user fatigue. These are key elements to build an innovation pump.

4 The innovation pump

4.1 The components

An innovation pump can be decomposed in five key elements:

- the user
- a subjectively evaluated problem to solve
- an interface to collect user's evaluations
- modelling and learning user preferences
- generate knowledgeable guesses of high quality solutions.

The innovation pump relies on the user to provide relative evaluations of the quality of the potential solutions to the problem at hand. Such assessment cannot be provided by any other means than user expertise and subjectivity. Thus, the innovation pump collects these evaluations – and other interactive aspects of user interaction – via the properly designed user interface. This interactive process between the user and the computer gives away information about what kind of solutions are considered high-quality candidates. This harvested information is the base to model and learn what are the elements that lead to high-quality solutions – the core of the innovation pump. Finally, once the learned user preferences are available, an innovation pump can exploit such preferences to generate high-quality solution candidates to be presented to the user for evaluation.

4.2 The challenge: user fatigue and frustration

Two key elements in the road to success to any interactive innovation pump is the user perception of the usefulness of the process. For instance, user fatigue is a critical element to produce high-quality solutions. Long times to first hit a high-quality solution may demand tedious and extended periods on the user focused attention. Some simple numbers can help to understand the boundaries of user attention. Let us assume that the user spends a minute to compare two candidate solutions. This means that the user can compare, no more than 120 solutions per hour. Moreover, assume that these evaluations are a repetitive process. If no high-quality solutions are quickly provided early on the interactive process, the user attention will drop and will start feeling fatigued to conduct and interactive process that is not providing an interesting answer (Takagi, 2001). Fatigue becomes the main reason of an early stop of the iGA process and, hence, leads to low-quality solutions. However it also has another by product: frustration. The more the user advances on the interactive process, if no high quality solutions are provided, the frustration for not achieving any useful solution in the elapsed time spent usually leads to early stops of the interactive search process.

Thus, it is key to the success of an interactive innovation pump to address the fatigue and frustration. As we have already mentioned, learning the user preferences about high-quality solutions, and using the gathered knowledge to generate educated guesses about potential high-quality solutions can alleviate the user fatigue – by reaching a good solution faster – and frustration – by showing for evaluation only promising solutions.

4.3 Human-computer collaboration for creative processes


The main goal of an innovation pump is to assist the user along the process of problem solving. When an engineer sits down to deal with a problem, besides any numerical modelling that can be made, intuition, experience, and background knowledge are powerful tools. Most of this knowledge is not easy verbalise or formalise sometimes, but it can be quickly accessed and recalled as suggested by Gladwell (2005). However, the user easily gives the necessary tips away every time he provides a solution evaluation. Thus, models of his preferences can be learned. Moreover, the learned models of preferences can be exploited to generate new high-quality candidate solutions – or *educated guesses*.

4.4 A minimalist approach

Active interactive GAs (Llorà et al., 2005) build an innovation pump to fuse human and computer efforts for problem solving. They collect the evaluations provided by the user – this can be easily generalised to groups of users – and builds models of the preferences. This subsection presents how the user evaluations can be collected and stored to allow the creation of user preferences.

In a tournament selection of size $s = 2$, a user is asked to provide an answer to the question of which of the two choices is better. The outcome of such question may be: the first shown, the second shown, or both are equal or the user was unable to decide. Let us assume the illustrative example where eight solutions need to be evaluated as presented in Figure 3.

Figure 3 Eight randomly chosen individuals from a population, are grouped in seven different tournaments 010111, 010100, 010101, 100001, 100000, 101010, 001000, 001110, 010111, 010101, 100000, 001000, 010111, 100000. The number beside each node simulates the objective function in the user mind



The tournament ordering presented in Figure 3 guarantees that the partial order introduced by the use revaluations produces a connected graph G . Such graph $G = \langle V, E \rangle$ represents the partial evaluation order representing the solutions as vertex in V , and the pair-wise comparison among individuals (greater, lesser, or equal) as edges in E . The partial ordering graph provided by the user may be undirected (equal evaluations are allowed), however, such graph can be easily turned into a directed graph as Figure 4 shows. The directed graph is obtained by replacing the equal (undirected edges) by the proper greater or lesser relations (directed edges), as Figure 4 shows.

Figure 4(a) Partial ordering graph provided by the comparisons provided by the user based on the tournaments of Figure 3





Figure 4(b) Equivalent partial order graph where equality relations have been replaced by the proper greater than relations



To create a synthetic fitness we propose a heuristic based on the partial ordering provided by user evaluations and the Pareto dominance concept (Pareto, 1896) of multi objective optimisation (Coello-Coello, 1998; Deb et al., 2000). A global ordering measure may be computed using a heuristic based on two dominance measures, δ and ϕ . Let us define $\delta(v)$ as 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 . Since the partial order is a directed graph, such mapping has an interesting property. If v appears more than once in a path (trial) of $\delta(v)$ or $\phi(v)$, then a cycle in such graph exists. Thus, due the greater and lesser relations, a contradiction on the user evaluations is identified. We will not discuss this issue further, leaving it as part of the further research. The work presented breaks such cycles by removing the oldest evaluation (edge) in the path thus breaking the cycle. Table 1 computes $\delta(v)$ and $\phi(v)$ given the graph presented in Figure 4(b).

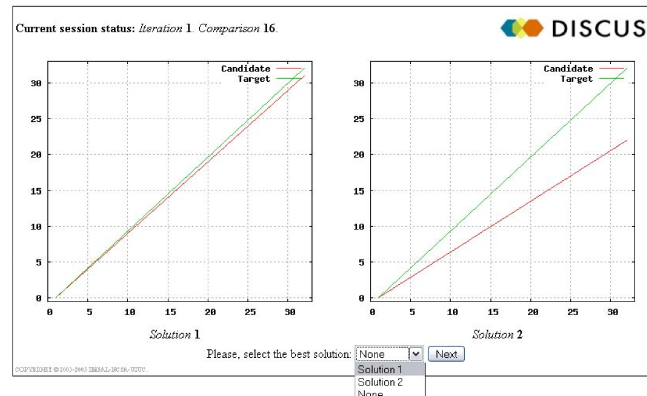
Table 1 Estimation of the global ranking based on the dominance measure. This data presented on the table uses the partial order presented on Figure 4(b). Given a vertex v , the number of dominated vertexes $\delta(v)$ and dominating vertexes is computed. Using these measures, the estimated fitness may be computed as $\hat{f}(v) = \delta(v) - \phi(v)$. The estimated ranking is obtained by sorting based on $\hat{f}(v)$

v	$f(v)$	$r(v)$	$\delta(v)$	$\phi(v)$	$\hat{f}(v)$	$\hat{r}(v)$
010111	4	1	5	0	5	1
010100	2	3	0	1	-1	4
010101	3	2	1	1	0	3
100001	2	3	0	2	-2	5
100000	1	4	0	3	-3	6
101010	3	2	2	0	2	2
001000	1	4	0	3	-3	6
001110	3	2	2	0	2	2

The estimated fitness of a given solution (vertex) v may be computed as $\hat{f}(v) = \alpha(v) - \phi(v)$. Intuitively, the more solutions a solution 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 ordering $\hat{r}(v)$ is obtained sorting by $\hat{f}(v)$, as shown in Table 1. The estimated ranking introduces some spurious relations inside common ranks. Once the global ordering is computed (estimated ranking $\hat{r}(v)$), such ordering may be used to train a ε -SVM. By optimising such a synthetic fitness we may obtain a look-a-head on candidate solutions to be evaluated by the user.

In our first validation experiments (Llorà et al., 2005) we created a simple web interface to test the performance of the *active* iGA proposed in Table 1. The interface was designed to minimise the interface bias providing a clear goal definition, a simple problem visualisation, and a clear relative comparison method to help maintain the user criteria. Figure 5 shows a snapshot of a real iGA session conducted using the proposed *active* iGA.

Figure 5 A simple interface for the *active* iGA test (see online version for colours)



4.5 Learning by interaction

Active interactive GAs rely on learning from the interaction with the user and anticipate what hypotheses the user may be interested via *educated guesses*. The idea is simple: mine the information provided by the user and use the obtained knowledge to guide the breeding process of new solutions in an educated manner. However, achieving such a goal requires to address several aspects of the interaction with the user that are beyond the scope of this paper and are introduced elsewhere (Llorà et al., 2005).

The algorithmic description of a iGA is summarised in Algorithm 1. The description presented by Llorà et al. (2005) assumed a theoretical framework where the decision variables involved on the interactive process were independent to simplify the theoretical analysis done. Such an assumption does not constrain the capability of a iGA to achieve linkage-learning capabilities (Goldberg, et al., 1989; Goldberg, 2002) provided proper learning and optimisation algorithms are chosen. As a result of the theoretical framework adopted, the initial a iGA used a ε -SVM (Vapnik, 1998, 1999; Cristianini and Shawe-Taylor, 2000; Shawe-Taylor and Cristianini, 2004) with a linear kernel as the model for the synthetic fitness and the Compact Genetic Algorithm (cGA) (Harik et al., 1998) to

optimise the learned synthetic fitness. Further details and motivation of these decisions may be found elsewhere (Llorà et al., 2005). The consistency measure proposed on this paper relies on the partial-ordering graph G which is incrementally built in lines 2, 3, 4, and 9 of the steps presented in Algorithm 1.

Algorithm 1 Algorithmic description of the aiGA as presented in Llorà et al. (2005)

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

4.6 The key: knowledgeable recombination

Building surrogate models of the synthetic fitness is a crucial milestone toward *educated guesses*. They can be used to compare the relative goodness of two solutions. Hence, we can use search methods to obtain *educated guesses* – promising solutions. This process can be regarded as a *knowledgeable recombination* of the user preferences in a quest to innovate and create new promising high-quality solutions. The *active* iGA (Llorà et al., 2005) uses this idea of searching among learned preferences.

The algorithmic description of aiGA is summarised in Table 1. Given the theoretical framework used so far (One Max and ε -SVM using a linear kernel), the compact GA is a suitable option to optimise the synthetic fitness. The cGA (Harik et al., 1998), is one of the simplest estimation of distribution algorithms (Pelikan et al., 2002; Larrañaga and Lozano, 2002). Similar to other EDAs, cGA replaces traditional variation operators of GAs by building a probabilistic model of promising solutions and sampling the model to generate new candidate solutions. The probabilistic model used to represent the population is a vector of probabilities, and therefore implicitly assumes each gene (or variable) to be independent of the other. Specifically, each element in the vector represents the proportion of ones (and consequently zeros) in each gene position. The probability vectors are used to guide further search by generating new candidate solutions variable by variable according to the frequency values.

The cGA consists of the following steps:

- 1 *Initialisation*. As in simple GAs, where the population is usually initialised with random individuals, in cGA we start with a probability vector where the probabilities are initially set to 0.5. However, other initialisation procedures can also be used in a straightforward manner.
- 2 *Model sampling*. We generate two candidate solutions by sampling the probability vector. The model sampling procedure is equivalent to uniform crossover in simple GAs.
- 3 *Evaluation*. The fitness or the quality-measure of the individuals is computed.

- 4 *Selection.* Like traditional genetic algorithms, cGA is a selectionist scheme, because only the better individual is permitted to influence the subsequent generation of candidate solutions. The key idea is that a “survival-of-the-fittest” mechanism is used to *bias* the generation of new individuals. We usually use tournament selection (Goldberg, 1989) in cGA.
- 5 *Probabilistic model updating.* After selection, the proportion of winning alleles is increased by $1/n$. Note that only the probabilities of those genes that are different between the two competitors are updated. That is,

$$p_i^{t+1} = \begin{cases} p_i^t + 1/n & \text{if } x_{w,i} \neq x_{c,i} \text{ and } x_{w,i} = 1 \\ p_i^t - 1/n & \text{if } x_{w,i} \neq x_{c,i} \text{ and } x_{w,i} = 0 \\ p_i^t & \text{Otherwise.} \end{cases}$$

where, $x_{w,i}$ is the i th gene of the winning chromosome, $x_{c,i}$ is the i th gene of the competing chromosome, p_i^t is the i th element of the probability vector – representing the proportion of i th gene being one – at generation t . This updating procedure of cGA is equivalent to the behaviour of a GA with a population size of n and steady-state binary tournament selection.

- 6 Repeat steps 2–5 until one or more termination criteria are met.

More details are available elsewhere (Harik et al., 1998, 1999). However it is important to note that the cGA is operationally equivalent to the order-one behaviour of simple genetic algorithm with steady state selection and uniform crossover (Harik et al., 1998). Therefore, the theory of simple GAs can be directly used in order to estimate the parameters and behaviour of the cGA.

The cGA provides a provability vector that – if not pushed till convergence – provides a bias joint distribution of the user preferences thanks to the optimisation of the synthetic fitness. Thus, sampling the final probability vector provides by cGA provide samples around the areas of learned areas of interest of the user. These samples inherently are the result of the recombination of solutions that contained different facets of the user’s preferences and, thus, they become the *educated guesses* to be shown for user evaluation.

5 Pumping in the real world

This section presents some of the results achieved using innovation pumps. It briefly reviews the results achieved after analysing the scalability of an iGAs. It also discusses the results of using a iGAs on a real-world problem: the weight tuning of a Text-To-Speech (TTS). Results show how, with the proper support for innovation and creativity, collaborative engineering provided speech engineers with a new and efficient methodology to tune TTS systems using subjective perception.

5.1 The approach

Before trying the active iGA on real problems, we conducted a detailed analysis of its behaviour using the interface presented in Figure 5. The first difference of that using a synthetic fitness function is the population size requirements. Figure 6(a) presents the population sizing of a simple GA and the one of the iGA. As it can be seen, the iGA requires a population size that, at least, grows linearly. Such requirement is the result of using a ϵ -SVM with a polynomial kernel which require at least as many training examples as dimensions (l in the iGA case). Moreover, the active iGA population is also constrained by the tree tournament structure; given a problem size l , the population size is forced to grow $2^{\lceil \log_2(l) \rceil}$.

Figure 6 Analysis of the results obtained using the *active* iGA propose when compare to a simple iGA. (a) population size; (b) convergence time; (c) function evaluations and (d) speed up (see online version for colours)

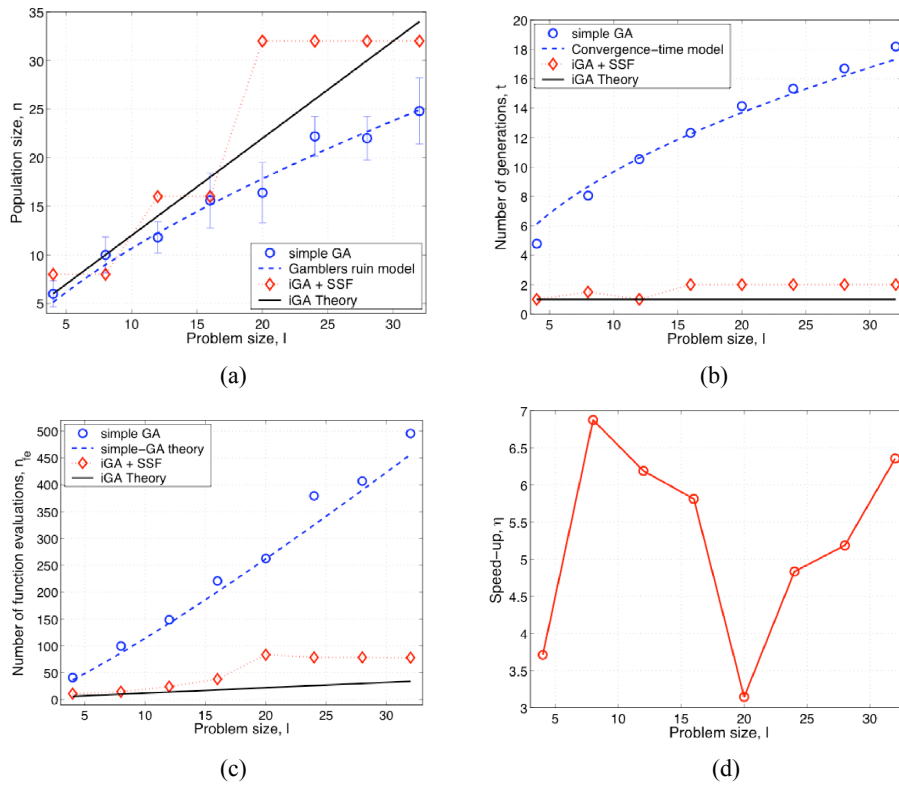


Figure 6(b) compares the convergence time of a active iGA to the a simple iGA. The theoretical convergence time of the active iGA with the proper population sizing should be constant. The empirical results show in Figure 6(b) support such assumption. Combining the population sizing and the convergence time, the number of functions evaluations of the active iGA should grow linear. However, due to the three structure of the tournament evaluation used, a staircase effect may be appreciated in Figure 6(c).

Finally, Figure 6(d) shows the speedup achieved using the active iGA respect to a simple iGA. The results show how with the active use of a simple low-cost high-error synthetic fitness function we were able to achieve speedups ranging from 3 up to 7 times. The instability of the speedup is the result of the constraints on the population sizing (See Figure 6(a)). However, being able to cut down the total number of evaluations on such ratios proved to be an effective method for combating the user fatigue.

5.2 The collaborative tuning of Text-To-Speech systems

The aim of any TTS system is the generation of synthetic speech from text. Human beings based on the perceived speech quality evaluate the performance of such systems. Hence, it is essential to somehow embed this subjective criterion into the tuning process of the TTS system for achieving highly natural synthetic speech. The corpus-based or *unit selection* TTS approach is one of the state-of-the-art techniques that try to reach this aim (Black and Tokuda, 2005). This method generates the synthetic speech signal by means of the selection and concatenation of recorded speech units. The tuning of the unit selection module is one of the most important processes in getting high quality synthetic speech (Black, 2002). The selection process is driven by a cost function (Hunt and Black, 1996), which is typically computed as the combination of several weighted subcosts. A key issue involves the accurate tuning of these weights, that is, mapping the user subjective preferences among candidate units – a complicated task (Hunt and Black, 1996; Lee et al., 2001). Several approaches have been proposed for weight training, distinguishing between

- hand-tuning (Coorman et al., 2000)
- machine-driven – purely objective methods (Hunt and Black, 1996; Meron and Hirose, 1999; Park et al., 2003) or perceptually optimised techniques (Lee et al., 2001; Peng et al., 2002; Toda et al., 2004).

Alías and Llorà introduced GAs for tackling the weight-tuning problem (Alías and Llorà, 2003). This technique overcame the restrictions of classic approaches (Hunt and Black, 1996; Meron and Hirose, 1999), attaining better results with a feasible computational effort. Nevertheless, this approach, as all the previous techniques, needs to face a key challenge: the reliable estimation of the subjective perception of the speech attributes (i.e., it is very difficult to define a solid perception mapping function). Thus, it is necessary to actually incorporate user preferences for accurately tuning the weights of the cost function. As a first step, we applied a simple interactive genetic algorithm for weight tuning, allowing an actual perception-guided adjustment (Alías et al., 2004). However, the conducted experiments evidenced two main problems: the tediousness of the process (user fatigue) and the complexity of maintaining a stable comparison criterion throughout the whole process (user consistency), which are weaknesses related to iGAs. Later, *active* iGAs (Llorà et al., 2005) showed that learning from user interaction and exploiting the learned knowledge to guide the process of collecting user evaluations can greatly reduce the number of evaluations required to achieve high-quality solutions. Moreover, as described in (Llorà et al., 2006b), the user consistency can be evaluated according to the number of cycles of the partial-ordering graphs.

The innovation pump modelled by a iGA was used to replace the original iGA and to fight the user's fatigue, frustration, and evaluation inconsistency. The usage of

partial-ordering graphs allowed quantifying the consistency of the solution-quality assessments provided by the users. A detailed description is beyond the scope of this paper can be found elsewhere (Llorà et al., 2006b). Results showed that a iGAs slashed in half the number of evaluations required achieving efficient subjectively tuned weights, reducing user fatigue during the tuning process. Moreover, the active iGAs also provided better user guidance, drastically boosting the user consistency along the tuning process when compared to a traditional iGA scheme.

6 Conclusions

This paper has reviewed how the pervasive expansion of computers and internet is changing the everyday work of engineers. Distributed and scalable environments for the integration of both human- and computer-generated knowledge are rising as critical commodities that are making the field of collaborative engineering a reality. Cybercollaboratories and cyberinfrastructure are positioning themselves as key players to guarantee the effectiveness of human-human and HMC, innovation, and creativity. We have also pointed out the key elements and mechanisms to put in place to support innovation and creativity on such distributed and heterogeneous environments. The DISCUS project has been briefly reviewed as one of the initial paradigms to support such requirements and being already used in the field of collaborative engineering.

We have also reviewed how the usage of tools of modern artificial and computational intelligence can create effective computational embodiments for creativity and innovation support. These tools help participants in collaborative creative processes to reflect and innovate using visual abstractions and representations of the main topics of discourse – to mention one of the possible modes of usage. In particular, we have focused on *innovation pumps*, computational embodiments that assist and guide engineers in problem-solving activities. The innovation pump learns the inherent preferences of the participants and exploits the learned knowledge to provide *educated guesses* of high-quality candidates for user evaluation. Results on real-world problems have shown how innovation pumps can provide substantial speedups to creative processes at the same time that they fight user fatigue, frustration, and evaluation inconsistency.

Acknowledgements

The work conducted in the DISCUS project could not have been possible without the help and friendship of the people in the IlliGAL, the ALG at NCSA, the Chance Discovery Consortium, and Hakuhodo, Inc. This work was sponsored by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF, under grant F49620-03-1-0129, the Technology Research, Education, and Commercialisation Center (TRECC), at University of Illinois at Urbana-Champaign, administered by the NCSA and funded by the Office of Naval Research under grant N00014-01-1-0175 and the National Science Foundation under grant IIS-02-09199. The US Government is authorised 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 Technology Research, Education, and Commercialisation Center, the Office of Naval Research, the National Science Foundation, or the US Government.

References

- Alías, F. and Llorà, X. (2003) 'Evolutionary weight tuning based on diphone pairs for unit selection speech synthesis', *EuroSpeech*, Geneva, Switzerland, Vol. 2, pp.1333–1336.
- Alías, F., Llorà, X., Iriondo, I., Sevillano, X., Formiga, L. and Socoró, J.C. (2004) 'Perception-guided and phonetic clustering weight tuning based on diphone pairs for unit selection TTS', *Proceedings of the 8th International Conference on Spoken Language Processing (ICSLP)*, Jeju Island, Korea, pp.1333–1336.
- Bacardit, J., Bernado-Mansilla, E., Butz, M., Kovacs, T., Llorà, X. and Takadama, K. (2008). *Learning Classifier Systems: 10th International Workshop, IWLCS 2006, Seattle, MA, USA, 8 July, 2006, and 11th International Workshop, IWLCS 2007, London, Revised Selected Papers*, Lecture Notes in Artificial Intelligence, Springer, Berlin.
- Black, A.W. (2002) 'Perfect synthesis for all of the people all of the time', *IEEE TTS Workshop 2002 (Keynote)*, Santa Monica, USA.
- Black, A.W. and Tokuda, K. (2005) 'Blizzard Challenge – 2005: evaluating corpus-based speech synthesis on common datasets', *Proceedings of the 9th International Conference on Speech Communication and Technology (InterSpeech)*, Lisboa, Portugal, pp.77–80.
- Caldwell, C. and Johnston, V.S. (1991) 'Tracking a criminal suspect through face-space with a genetic algorithm', *Proceedings of the Fourth International Conference on Genetic Algorithms*, Morgan Kaufmann, pp.416–421.
- Coello-Coello, C. (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, December.
- Coorman, G., Fackrell, J., Rutten, P. and Van Coile, B. (2000) 'Segment selection in the L&H real speak laboratory TTS system', *ICSLP*, Beijing, China, Vol. 2, pp.395–398.
- Cristianini, N. and Shawe-Taylor, J. (2000) *An Introduction to Support Vector Machines*, Cambridge Press, Cambridge, MA.
- 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, Kampur.
- Gladwell, M. (2005) *Blink: The Power of Thinking without Thinking*, Brown Little, Boston.
- Goldberg, D.E. (1983) *Computer-aided Gas Pipeline Operation using GAs and Rule Learning*, Dissertation Abstracts International, Doctoral Dissertation, University of Michigan, Vol. 44, No. 10, p.3174B, Ann Harbor, MI.
- Goldberg, D.E. (1989) *Gas in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, MA.
- Goldberg, D.E. (1994) 'Genetic and evolutionary algorithms come of age', *Communications of the ACM*, Vol. 37, No. 3, pp.113–119.
- Goldberg, D.E. (2002) *The Design of Innovation: Lessons from and for Competent Genetic Algorithms*, Kluwer Academic Publisher, Boston, MA.
- Goldberg, D.E., Deb, K., Kargupta, H. and Harik, G. (1993) 'Rapid, accurate optimization of difficult problems using fast messy genetic algorithms', *Proceedings of the Fifth International Conference on Genetic Algorithms (ICGA'93)*, Morgan Kaufmann, Saint Louis, MO, pp.56–64.
- Goldberg, D.E., Korb, B. and Deb, K. (1989) 'Messy genetic algorithms: motivation, analysis, and first results', *Complex Systems*, Vol. 3, No. 5, pp.493–530.

- Goldberg, D.E., Welge, M. and Llorà, X. (2003) *DISCUS: Distributed Innovation and Scalable Collaboration in Uncertain Settings*, IlliGAL Report No. 2003017, University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory, IL, Urbana.
- Harik, G., Cantú-Paz, E., Goldberg, D.E. and Miller, B.L. (1999) 'The gambler's ruin problem, genetic algorithms, and the sizing of populations', *Evolutionary Computation*, Also IlliGAL Report No. 96004, Vol. 7, No. 3, pp.231–253.
- Harik, G., Lobo, F. and Goldberg, D.E. (1998) 'The compact genetic algorithm', *Proceedings of the IEEE International Conference on Evolutionary Computation*, Also IlliGAL Report No. 97006, pp.523–528.
- Hunt, A. and Black, A.W. (1996) 'Unit selection in a concatenative speech synthesis system using a large speech database', *ICASSP*, Atlanta, USA, Vol. 1, pp.373–376.
- Kosorukoff, A. and Goldberg, D.E. (2002) 'Evolutionary computation as a form of organization', *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, Morgan Kaufmann, pp.965–972.
- Kovacs, T., Llorà, X., Takadama, K., Lanzi, P.L., Stolzman, W. and Wilson, S.W. (2007) *Learning Classifier Systems: International Workshops, IWLCS 2003–2005*, Revised Selected Papers, Lecture Notes in Artificial Intelligence, Springer, Berlin.
- Larrañaga, P. and Lozano, J.A. (Eds.) (2002) *Estimation of Distribution Algorithms*, Kluwer Academic Publishers, Boston, MA.
- Lee, M., Lopresti, D.P. and Olive, J.P. (2001) 'A text-to-speech platform for variable length optimal unit searching using perceptual cost functions', *Fourth ISCA Workshop on Speech Synthesis*, Perthshire, Scotland, pp.75–80.
- Llorà, X., Alías, F., Formiga, L., Sastry, K. and Goldberg, D.E. (2006a) 'Efficient interactive weight tuning for TTS synthesis: reducing user fatigue by improving user consistency', *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Toulouse, France (Also IlliGAL Report No. 2005022), Vol. 1, May, pp.865–868.
- Llorà, X., Goldberg, D.E., Ohsawa, Y., Matsumura, N., Washida, Y., Tamura, H., Yoshikawa, M., Welge, M., Auvil, L., Searsmith, D., Ohnishi, K. and Chao, C.J. (2006b) 'New mathematics and natural computation', *World Scientific*, Vol. 2, No. 1, pp.85–100.
- Llorà, X., Goldberg, D.E., Ohsawa, Y., Ohnishi, K., Tamura, H., Washida, Y. and Yoshikawa, M. (2004a) 'Chances and marketing: online conversation analysis for creative scenarios', *First European Workshop on Chance Discovery (EWCD)*, pp.141–152.
- Llorà, X., Matsumura, N., Goldberg, D.E., Ohsawa, Y., Ohnishi, K. and Gonzales, K. (2004b) 'Discovering chance scenarios using small-world key graphs and evolutionary computation', In *First European Workshop on Chance Discovery (EWCD)*, pp.51–62.
- Llorà, X., Ohnishi, K., Chen, Y.P., Goldberg, D.E. and Welge, M. (2004c) 'Enhanced innovation: a fusion of chance discovery and evolutionary computation to foster creative processes and decision making', *Genetic and Evolutionary Computation Conference (GECCO)*, also as IlliGAL Report No. 2004012, pp.1314, 1315.
- 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', *Proceedings of Genetic and Evolutionary Computation Conference*, GECCO, Also IlliGAL Report No. 2005009, pp.1363–1371.
- Matsumura, N. (2003) *Chance Discovery, Chapter Topic Diffusion in a Community*, Springer, Berlin, pp.84–97.
- Matsumura, N., Goldberg, D.E. and Llorà, X. (2005) 'Mining social networks in message boards. In the artificial intelligence and simulation of behavior (AISB)', *Symposium on Conversational Informatics for Supporting Social Intelligence and Interaction*, Also IlliGAL Report No. 2005001.
- Meron, Y. and Hirose, K. (1999) 'Efficient weight training for selection based synthesis', *Euro Speech*, Budapest, Hungary, Vol. 5, pp.2319–2322.

- Ohsawa, Y. (2002) 'Chance discoveries for making decisions in complex real world', *New Generation Computing*, Vol. 20, No. 2, pp.143–164.
- Pareto, V. (1986) *Cours d'Economie Politique*, Vols. I and II, F. Rouge, Lausanne.
- Park, S.S., Kim, C.K. and Kim, N.S. (2003) 'Discriminative weight training for unit-selection based speech synthesis', *Proceedings of the 8th European Conference on Speech Communication and Technology (Euro Speech)*, Geneva, Switzerland, Vol. 1, pp.281–284.
- Pelikan, M. and Goldberg, D.E. (2001) 'Hierarchical Bayesian optimization algorithm = Bayesian optimization algorithm + niching + local structures', *Genetic and Evolutionary Computation Conference (GECCO)*, San Francisco, California, Morgan Kaufmann, USA, pp.217–221.
- Pelikan, M., Lobo, F. and Goldberg, D.E. (2002) 'A survey of optimization by building and using probabilistic models', *Computational Optimization and Applications*, Also IlliGAL Report No. 99018, Vol. 21, pp.5–20.
- Peng, H., Zhao, Y. and Chu, M. (2002) 'Perpetually optimizing the cost function for unit selection in a TTS system with one single run of MOS evaluation', *Proceedings of ICSLP, Denver, USA*, International Speech Communication Association Press, Bonn.
- Shawe-Taylor, J. and Cristianini, N. (2004) *Kernel Methods for Pattern Analysis*, Cambridge Press, Cambridge, MA.
- Simon, H.A. (1969) *The Sciences of the Artificial*, MIT Press, Cambridge, MA.
- Takagi, H. (2001) 'Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation', *Proceedings of the IEEE*, Vol. 89, No. 9, pp.1275–1296.
- Toda, T., Kawai, H. and Tsuzaki, M. (2004) 'Optimizing sub-cost functions for segment selection based on perceptual evaluations in concatenative speech synthesis', *Proceedings of ICASSP*, Montreal, Canada, pp.657–660.
- Vapnik, V.N. (1998) *Statistical Learning Theory*, John Wiley & Sons, Inc., New York, NY.
- Vapnik, V.N. (1999) *The Nature of Statistical Learning Theory*, Springer-Verlag, Inc., New York.
- Welge, M., Auvil, L., Shirk, A., Bushell, C., Bajcsy, P., Cai, D., Redman, T., Clutter, D. Aydt, R. and Tchong, D. (2003) *Data to Knowledge (D2K) Technical Report*, Technical Report Automated Learning Group, National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign, Urbana, IL.

Notes

¹<http://www-discus.ge.uiuc.edu/>

²An example of a freely available human-based GA may be seen at www.3form.com in the Free Knowledge Exchange (FKX).

³<http://alg.ncsa.uiuc.edu/>