Towards Evolutionary Story Generation

Andrés Gómez de Silva Garza
Computer Engineering Department
Instituto Tecnológico Autónomo de México (ITAM)
Río Hondo #1, Colonia Progreso Tizapán, 01080—México, D.F. México
agomez@itam.mx

Rafael Pérez y Pérez
Departamento de Tecnologías de la Información, División de Ciencias de la Comunicación y Diseño
Universidad Autónoma Metropolitana, Cuajimalpa
México, D.F. México
rperez@correo.cua.uam.mx

Abstract

In this paper we describe on-going work on combining two existing models of computational creativity. The GENCAD model proposes the use of an evolutionary algorithm (EA) that uses a population of exemplars as a starting point for its search, unlike traditional EA's, which use a randomly-generated initial population. The EA, operating on this population, is then used to generate new potentially creative solutions. GENCAD has been instantiated in the domains of structural design of tall buildings and feng shui-compliant residential floor-plan design. The MEXICA model also begins with a set of exemplars as a starting point, but it analyzes these exemplars based on a domain theory. The general theory that is obtained from analyzing the set of exemplars is then used to guide the generation of new solutions. MEXICA has been instantiated in the domain of plot generation for stories involving themes, characters and locations from the Mexica culture of ancient Mexico. In the hybrid model we propose in this paper, we combine the two models to generate plots for stories of the same sort that MEXICA generates, but using GENCAD's process model to do so.

Introduction

We have begun work on combining two existing models of computational creativity which were developed independently. The purpose of this combination is to produce a hybrid model that maximizes the advantages and minimizes the disadvantages of both original models.

The first of the original models we are working with is GENCAD (Gómez de Silva Garza 2000). The generative module in GENCAD takes a set of pre-existing exemplars of the type of thing that we would like to create and interprets it as the initial population of an evolutionary algorithm (EA) (Mitchell 1998). The EA's genetic operators are then used to generate complete new potential solutions (new examples of the type of thing that we would like to create). The EA's evaluation module uses domain and common-sense knowledge to assign a fitness value to each of these new solutions that serves to rank the old and new solutions so that only the best solutions survive across evolutionary generations. Convergence of the EA occurs when one of the new potential solutions is determined to be of sufficient quality according to both the EA's fitness function and whatever initial problem requirements the user may have specified. This process model has been instantiated in two domains: the structural design of tall buildings (Gómez de Silva Garza and Maher 1998), and the design of residential floor plans that follow the principles of feng shui (Gómez de Silva Garza and Maher 1999). In these two instantiations the set of exemplars that is used as the initial population of the EA results from an earlier phase of the process model which implements the case retrieval stage of a case-based reasoner, whereas the EA implements the case adaptation stage (Leake 1996).

The second of the original models we are working with is MEXICA (Pérez y Pérez 1999). The generative module in MEXICA takes a set of pre-existing exemplars of the type of thing that we would like to create and analyzes it according to a domain theory. The domain in which MEXICA has been instantiated is the generation of plots for stories involving themes, characters, and locations from the Mexica culture of ancient Mexico (Pérez y Pérez and Sharples 2001; Pérez y Pérez 2007), though some initial work has been done on applying the model to image layout design (Pérez y Pérez et al. 2012). The theory used by MEXICA for this domain is based on an analysis of the emotional links between characters and the flow (increase and decrease) of tensions in a story as actions take place in the story. As a result of analyzing the set of exemplars according to this domain theory, MEXICA produces an abstract description of the entire set. Given an initial action, it then starts to add more and more actions using the abstract description as a set of guidelines that ensure coherence, thus eventually producing a complete story piecemeal.

As can be seen, there are similarities between GENCAD and MEXICA, yet there are also quite a few, sometimes subtle, differences. In the next section of the paper we discuss these issues further, as well as the advantages and disadvantages of the two process models from the point of view of computational creativity. In the section after that we describe our hybrid model and some of its characteristics. Finally, in the last section we provide some results, lessons, and observations from our preliminary experiments with our hybrid model.

Comparing and Contrasting GENCAD and MEXICA

In this section we compare and contrast the evolutionary approach used by GENCAD with the theory-based approach used by MEXICA for the generation of solutions to computational creativity problems. This is done by analyzing some of the characteristics of the two approaches and their relative advantages and disadvantages.

One of the characteristics of evolutionary approaches is that the way in which they generate new potential solutions is generally syntactic rather than semantic. Existing genotypes are tweaked (by the mutation operator) or split and spliced in order to combine their characteristics (by the crossover operator) without any prior analysis of whether the results will "make sense" or not, or of the meaning of the genotypes. This analysis is left to the EA's evaluation module later on in the process. This means that the generative module is generally not biased or guided by any domain knowledge, thus increasing the potential for interesting, unexpected features in the generated solutions, an important characteristic in creativity (Grace and Maher 2014).

Another characteristic of evolutionary approaches is that many of the decisions in the generative module are made at random, such as, in the case of the crossover operator, which genotypes will be combined or where exactly they will be split (before splicing the resulting pieces to produce the resulting new genotypes). Thus even if the same algorithm is run again on the same initial population, the results are not likely to be the same as in previous runs. This unpredictability is another potential source of unexpectedness in the generated solutions. This characteristic also implies that, if for some reason convergence isn't reached during one attempt to process a given initial population, the attempt can be abandoned and a new attempt initiated, with the possibility that the new attempt will converge, thus providing the approach with more flexibility than traditional algorithms possess.

On the other hand, there are disadvantages to evolutionary approaches, which include the following. First, even if convergence is reached (that is, even if eventually a solution that is "good enough" is produced by the evolutionary process), most of the time many bad quality potential solutions may have had to be generated, through a large number of evolutionary generations, before convergence. In addition, even if the capability to "give up" (in order to re-start the evolutionary process to try again) is programmed into the EA, this usually has to be done after a large number of evolutionary generations in order to take into account the slow speed of evolution. In other words, EA's are generally not very efficient.

One of the characteristics of theory-based approaches is that the solutions that are generated are guaranteed from the first to "make sense" (unless the theory is deficient in some way, e.g., incorrect or incomplete). Thus, finding a solution that is "good enough" does not require wasting time on slowly discarding many more defective solutions that were also generated, which is what happens in an EA.

On the other hand, the solutions that are generated are always based on the theory, and by definition will never contain features that go beyond that theory. The constraints imposed by the theory may be too rigid to permit that spark of interestingness or unexpectedness that can be so important in creativity.

Further discussion of these issues involving theory-based approaches is included in the following section.

Hybrid Model Combining the Evolutionary- and Theory-Based Approaches

In order to combine the advantages of the evolutionary- and theory-based approaches to generating solutions for computational creativity systems, we have produced a hybrid model which we describe in this section. We are still in the process of instantiating this model in the domain of story generation.

Our hybrid model, like both GENCAD and MEXICA, begins with a set of exemplars. Following GENCAD's process model, these exemplars are treated as the initial population of an EA whose genetic operators are then used to generate new potential solutions. In our hybrid model the EA's evaluation module is implemented using a looser version of MEXICA's theory combined with common-sense constraints. Thus, some aspects of MEXICA's domain theory are used to guide the generative process and filter out the more deficient solutions, but the rigidity imposed by exclusively following the constraints imposed by the original theory when generating new solutions is counteracted by the flexibility introduced by the genetic operators.

Specifically for the domain of story generation, assuming that MEXICA's standard set of 7 pre-existing stories is used, part of the general, abstract description (theory) it would come up with after analyzing these 7 stories would be that if a character A likes a character B, and B likes another character C, then A becoming jealous of C is a possible next action to introduce to the story being generated. This description arises from the fact that this type of situation (sequence of story actions) occurs in the pre-existing stories. In fact, unless there are multiple other possible next actions that can be introduced at a given point in MEXICA's creation of a story in which this observation is relevant, A becoming jealous of C will be the next action that will be introduced. In our new hybrid model, a story in which A becomes jealous of C shortly after it is stated that A likes B and B likes C will be assigned a higher fitness value than one in which it happens much afterward, and an even higher fitness value than one in which it doesn't happen at all. But these other possibilities are still present, thus increasing the variety in the structure of the new potential stories that can be generated.

The hybrid system's evaluation module also incorporates knowledge about the flow of dramatic tensions in "good" stories (the tension usually increases steadily up to a certain point, near the end, when there is usually a denouement during which all of the accumulated conflict and tension is resolved) as well as other aspects of MEXICA's domain theory. However, it turned out to be necessary to implement additional common-sense domain constraints in the evaluation module that never had to be represented explicitly in the original instantiations of MEXICA.

For instance, the flexibility of the genetic operators implies that, after several evolutionary generations, new stories that have "incestuous" ancestry may be created. Thus, if AB is a story created in generation 1 whose direct ancestors are A and B, then in generation 2 a new story ABA may be created whose direct ancestors are AB and A, thus containing some genetic material directly inherited from A and some genetic material indirectly inherited from A through AB. This may result in stories in which the sequence of actions is, for instance:

```
M H K L K M P
```

In other domains the potential repetitiveness inside a genotype (M and K appear twice in the example sequence given above) may not be important, or may even be desirable—it all depends on the interpretation of the contents of a genotype and on the application domain. However, in story generation the quality of a story is diminished if the author constantly repeats things that have already been stated instead of moving forward with new actions/events. Thus, our hybrid system takes this potential repetition into account when assigning a fitness value to the stories it generates. Further work is still being performed in order to identify which additional such common-sense constraints may be necessary, and in order to implement them in the fitness evaluation module.

Discussion, Results, and Lessons Learned

We have presented a hybrid model of computational creativity that combines aspects of two previously-existing models. The hybrid model uses an evolutionary algorithm (EA) for the generation of solutions, and implements the EA's evaluation module based on a domain theory arising from an analysis of exemplars of good solutions in the application domain. We have instantiated this hybrid model in the domain of story generation in order to test and refine our ideas.

Some work has been done in the past on using EA's for linguistic creativity, but has focused on sentence (Vrajitoru 2003) or poetry (Manurung 2003) generation, rather than story (plot) generation. More similar to our work is (McIntyre and Lapata 2010), though unlike us they do not avoid the use of domain knowledge in the generation module of the EA.

While our work is still preliminary, one of the results we have been able to obtain from this research is to be able to state explicitly the advantages and disadvantages of the original models by comparing and contrasting them. This analysis is what led to our proposal for the hybrid model, which tries to maximize the combined advantages and minimize the combined disadvantages of the original models.

Acknowledgements

This work has been partially supported by Asociación Mexicana de Cultura, A.C. and by the National Council of Science and Technology in México (CONACYT), project number: 181561.

References

Gómez de Silva Garza, A. 2000. *An Evolutionary Approach to Design Case Adaptation*. Ph.D. Dissertation, The University of Sydney, Australia.

Gómez de Silva Garza, A., and Maher, M.L. 1998. A Knowledge-Lean Structural Engineering Design Expert System. *Proceedings of the Fourth World Congress on Expert Systems*, Mexico City, Mexico. 178-185.

Gómez de Silva Garza, A., and Maher, M.L. 1999. Evolving Design Layout Cases to Satisfy Feng Shui Constraints. *Proceedings of the Fourth Conference on Computer Aided Architectural Design in Asia (CAADRIA-99)*, Shanghai, People's Republic of China. 115-124.

Grace, K., and Maher, M.L. 2014 (to appear). What to Expect When You're Expecting: The Role of Unexpectedness in Computationally Evaluating Creativity. *Proceedings of the Fifth International Conference on Computational Creativity (ICCC '14)*, Ljubljana, Slovenia.

Leake, D.B. 1996. *Case-Based Reasoning: Experiences, Lessons, & Future Directions*. Menlo Park, California & Cambridge, Massachusetts: AAAI Press & The MIT Press.

Manurung, H.M. 2003. *An Evolutionary Algorithm Approach to Poetry Generation*. Ph.D. Dissertation, The University of Edinburgh, Scotland.

McIntyre, N. and Lapata, M. 2010. Plot Induction and Evolutionary Search for Story Generation. *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, Uppsala, Sweden. 1562-1572.

Mitchell, M. 1998. *An Introduction to Genetic Algorithms (Complex Adaptive Systems Series)*. Cambridge, Massachusetts: MIT Press.

Pérez y Pérez, R. 1999. *MEXICA: A Computer Model of Creativity in Writing*. Ph.D. Dissertation, University of Sussex, United Kingdom.

Pérez y Pérez, R. 2007. Employing Emotions to Drive Plot Generation in a Computer-Based Storyteller. *Cognitive Systems Research* 8(2): 89-109.

Pérez y Pérez, R.; Morales, N.; and Rodríguez, L. 2012. Illustrating a Computer Generated Narrative. *Proceedings of the Third International Conference on Computational Creativity (ICCC '12)*, Dublin, Ireland. 103-110.

Pérez y Pérez, R., and Sharples, M. 2001. MEXICA: A Computer Model of a Cognitive Account of Creative Writing. *Journal of Experimental and Theoretical Artificial Intelligence* 13(2). 119-139.

Vrajitoru, D. 2003. Evolutionary Sentence Building for Chatterbots. *Late Breaking Papers, Proceedings of the Genetic and Evolutionary Computation Conference (GECCO '03)*, Chicago, Illinois. 315-321.