2014 IEEE International Conference on Data Mining Workshop

Commonsense Knowledge as the Glue in a Hybrid Model of Computational Creativity

Andrés Gómez de Silva Garza
Computer Engineering Department,
Instituto Tecnológico Autónomo de México (ITAM),
México, D.F., México
agomez@itam.mx

Erik Cambria
School of Computer Engineering,
Nanyang Technological University (NTU),
Singapore
cambria@ntu.edu.sg

Rafael Pérez y Pérez
Departamento de Tecnologías de la Información, Universidad Autónoma Metropolitana (UAM), Cuajimalpa,
México, D.F., México
rperez@correo.cua.uam.mx

***Abstract*—In this paper, we describe on-going work on the ensemble of two models for computational creativity using commonsense knowledge. The GENCAD model of computational creativity proposes the use of an evolutionary algorithm (EA) that employs a population of exemplars as a starting point for its search, unlike traditional EAs, which employ 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 of computational creativity also begins with a set of exemplars as a starting point, but it analyzes these exemplars based on a domain theory. The generalized model 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 combine the two original models to generate plots for stories of the same sort that MEXICA generates, but using GENCAD's process model to do so. To confirm that the generated plots make sense, the hybrid model integrates commonsense knowledge into its evaluation module.**

Keywords—commonsense knowledge; computational creativity; plot analysis and generation

I. 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 [5]. 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) [13]. The EA's genetic operators are then used to generate complete new potential solutions (i.e., new examples of the type of thing that we would like to create).

The EA's evaluation module uses domain 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, whether old or new, 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 [6], and the design of residential floor plans that follow the principles of feng shui [7]. 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. Viewed from the point of view of case-based reasoning, the EA implements the case adaptation stage [10].

The second of the original models we are working with is MEXICA [14]. 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 relevant to the Mexica culture of ancient Mexico [15, 17], though some initial work has been done on applying the model to image layout design [16]. 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.

This can be viewed as opinion mining within a set of exemplars of stories, resulting in one generalized representation of the relevant aspects of the entire set. Given an initial action, MEXICA then starts to add more and more actions using the abstract description as a set of guidelines that attempt to ensure coherence, thus eventually producing a complete story piecemeal.

978-1-4799-4274-9/14 $31.00 © 2014 IEEE
DOI 10.1109/ICDMW.2014.178

As can be seen, there are similarities between GENCAD and MEXICA, yet there are also quite a few, sometimes subtle, differences. In Section II, 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 Section III, we describe our hybrid model and some of its characteristics, as well as the role that commonsense knowledge plays and its impact in the model's performance. Finally, in the last section we provide some results, lessons, and observations from our preliminary experiments with our hybrid model.

II. 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 original 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 [8].

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 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 this new attempt will converge, thus providing the approach with more flexibility than traditional deterministic 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, and most of them discarded, before convergence. In addition, even if the capability to "give up" (in order to re-start the evolutionary process to try again, thus recovering from exploring apparently fruitless search paths) 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, EAs are generally not very efficient: a large amount of computer memory and processor time has to be invested in producing viable results.

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. This is a result of the fact that theory-based approaches use the semantics of a domain to guide the production of new solutions.

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.

III. 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 commonsense constraints [1]. We will discuss what we mean with “a looser version” a bit further on in the paper, in the context of an example from the domain of story generation. 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. A balance is struck between syntactic generation and semantic analysis/evaluation that combines the advantages of the original models but improves on some of the disadvantages.

Let us ground these observations using the domain of story generation. Normally the instantiation of MEXICA in this domain has performed its opinion mining on a set of seven pre-existing stories (using a formal representation for them, rather than their natural language representation, for expediency). Let us assume that this standard set of seven pre-existing stories is analyzed. Then part of the general, abstract description (theory) the system would come up with would be the "rule" 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 into the story being generated (remember that MEXICA constructs new stories piecemeal). This general "rule" 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, we can guarantee that A becoming jealous of C will be the next action that will be introduced. We can successfully predict that this will happen even before executing the system.

In order to aid such emotion recognition and story understanding dynamics, we plan to integrate sentic computing tools and techniques [2] into the model as a means for opinon mining and an aid for intention awareness [9] (see Fig. 1). By allowing sentiments to flow from concept to concept based on the dependency relation between clauses [18], sentic computing achieves a better understanding of the contextual role of concepts in each sentence and, hence, outperforms state-of-the-art statistical methods that simply rely on the syntactical representation of text. The sentence “iPhone5 is expensive but nice”, for example, is equal to “iPhone5 is nice but expensive” from a bag-of-words perspective. However, the two sentences bear opposite polarity: the former is positive as the user seems to be willing to make the effort to buy the product despite its high price, the latter is negative as the user complains about the price of iPhone5 although he/she likes it.

In sentic computing, whose term derives from the Latin sentire (root of words such as sentiment and sentience) and sensus (intended both as capability of feeling and as commonsense), the analysis of natural language is based on linguistic patterns and commonsense reasoning tools, which enable the analysis of text not only at document-, page- or paragraph-level, but also at sentence-, clause-, and concept-level. In particular, sentic computing involves the use of AI and Semantic Web techniques, for knowledge representation and inference; mathematics, for carrying out tasks such as graph mining and multi-dimensionality reduction; linguistics, for discourse analysis and pragmatics; psychology, for cognitive and affective modeling; sociology, for understanding social network dynamics and social influence; finally ethics, for understanding related issues about the nature of mind and the creation of emotional machines.

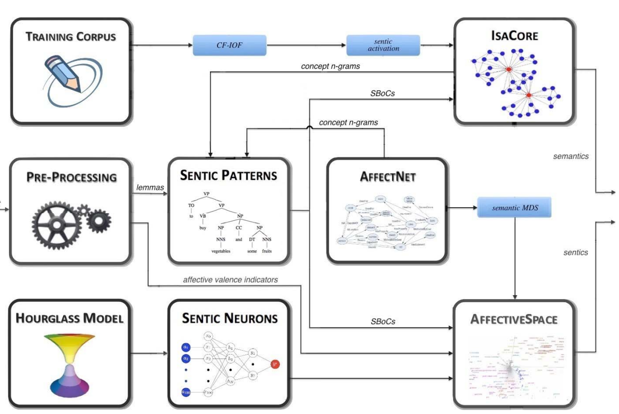

Figure 1. Sentic Computing schema

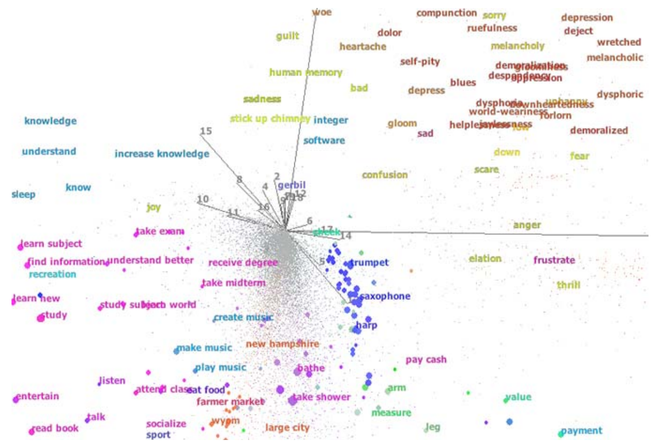

Figure 2. A sketch of SenticSpace

In particular, we want to exploit SenticSpace [3, 19] to assess emotions associated with actions, events, and narratives. SenticSpace (see Fig. 2) is a new framework for affective common-sense reasoning that extends WordNet-Affect and SenticNet by providing both emotion labels and polarity scores for a large set of natural language concepts. The framework is built by means of fuzzy c-means clustering and support-vector-machine classification, and takes into account different similarity measures, including point-wise mutual information and emotional affinity.

In our new hybrid model, remember that the EA generates several complete stories at a time. In this hybrid model, a generated 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 a story in which it doesn't happen at all. But these other possibilities are still present, thus increasing the variety in the structure and the difficulty in predicting the actual contents of the new potential stories that can be generated. This is what we mean when we say that the evaluation module of the hybrid system implements a looser version of MEXICA’s domain theory. It is looser in that it is less rigid about the variables/features it measures, not in the sense that it measures less variables/features. It is less rigid because it assigns a global fitness value to each generated story based on measuring multiple variables/features at the same time, and the global fitness value is obtained by combining all of the measurements (it is a linear combination, giving the same weight/importance to each measurement).

Thus, even if some aspects of a story are determined to be “not so good” by the evaluation module, if there are enough other aspects that are considered favorably, the story as a whole may end up being assigned a relatively high fitness value compared to the rest of the new stories. In the original instantiation of MEXICA it would simply not have been possible for new stories to end up having any aspect that is “not so good” (unless the original exemplar stories on purpose contain this type of feature).

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 commonsense domain constraints in the evaluation module that never had to be represented explicitly in the original instantiation of MEXICA. This will be aided by leveraging on affective commonsense knowledge and by adopting the Hourglass model (Fig. 3), a biologically-inspired and psychologically-motivated model for the representation and the analysis of human emotions [4].

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 a story 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. Our hybrid system now 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 commonsense constraints may be necessary, and in order to implement them in the fitness evaluation module. This is done by running the system, observing the resulting stories and their characteristics, and determining which features they contain that should really have been penalized by the hybrid system's evaluation module when assigning a fitness value. This information is then used to iteratively adjust and improve the evaluation module, and we will stop the iterations when we are happy with the overall quality of most of the stories that the hybrid system produces. This gradual improvement to the system is being done manually, so it is too early in the process to be able to present an evaluation of the (final) system's performance at this time.

In the original MEXICA this type of commonsense constraint never had to be represented explicitly because, for instance, none of the pre-existing stories analyzed by the system contained repeated actions. Even if they had, in the original MEXICA new stories are based on the semantics of pre-existing stories rather than their structure; so it never would have occurred to the system to repeat actions in a new story that it was generating. The hybrid system needs to explicitly implement these constraints because of the flexibility that the EA introduces into the types of stories that can be generated.

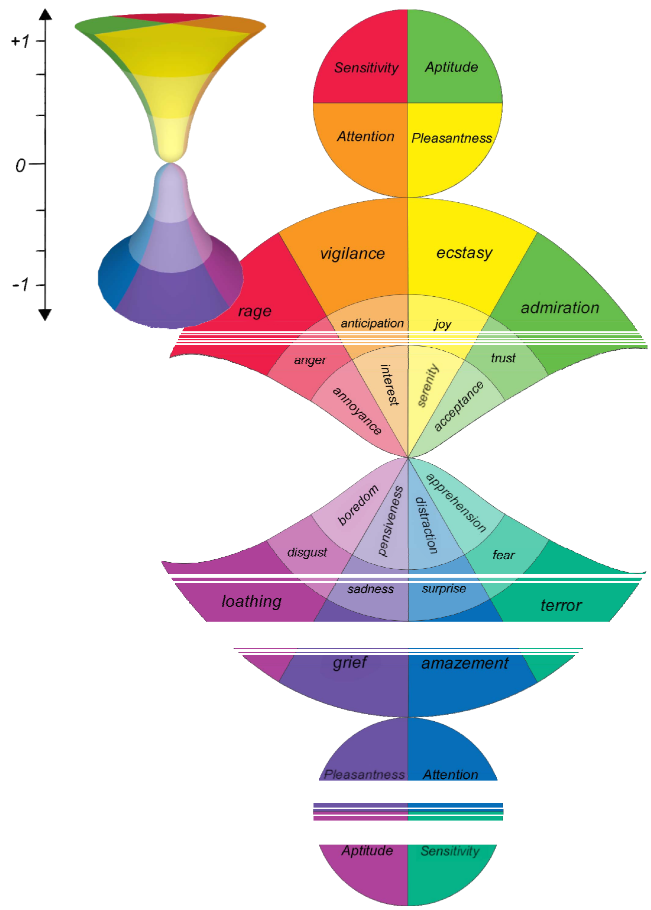

Figure 3. The Hourglass model

IV. 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 opinion mining within a set 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 EAs for linguistic creativity, but has focused on sentence [20] or poetry [11] generation, rather than story (plot) generation. More similar to our work is [12], though unlike us the participants in that project do not avoid the use of domain knowledge in the generative module of the EA, so the EA they use is not "pure."

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.

The other result we have obtained is the observation that commonsense knowledge can act as the glue that bridges the syntactic and semantic approaches to solution generation in computational creativity in order to obtain a successful hybrid model.

Acknowledgment

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

[1] E. Cambria, A. Hussain, C. Havasi, and C. Eckl. Common sense computing: From the society of mind to digital intuition and beyond. In: LNCS, vol. 5707, pp. 252-259, Springer, 2009.

[2] E. Cambria, A. Hussain, C. Havasi, and C. Eckl. Sentic computing: Exploitation of common sense for the development of emotion-sensitive systems. In: LNCS, vol. 5967, pp. 148-156, Springer, 2010.

[3] E. Cambria, A. Hussain, C. Havasi, and C. Eckl. SenticSpace: Visualizing opinions and sentiments in a multi-dimensional vector space. In: LNAI, vol. 6279, pp. 385–393, Springer, 2010.

[4] E. Cambria, T. Mazzocco, A. Hussain, and C. Eckl. Sentic medoids: Organizing affective common sense knowledge in a multi-dimensional vector space. In: LNCS, vol. 6677, pp. 601-610, Springer, 2011.

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

[6] A. Gómez de Silva Garza and M.L. Maher, "A knowledge-lean structural engineering design expert system," in Proceedings of the Fourth World Congress on Expert Systems, Mexico City, Mexico, pp. 178-185, 1998.

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

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

[9] N. Howard and E. Cambria. Intention awareness: Improving upon situation awareness in human-centric environments. Human-centric Computing and Information Sciences 3(9), 2013.

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

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

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

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

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

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

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

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

[18] S. Poria, E. Cambria, G. Winterstein, and G.-B. Huang. Sentic patterns: Dependency-based rules for concept-level sentiment analysis. Knowledge-Based Systems 69, pp. 45-63, 2014.

[19] S. Poria, A. Gelbukh, E. Cambria, A. Hussain, and G.-B. Huang. EmoSenticSpace: A novel framework for affective common-sense reasoning. Knowledge-Based Systems 69, pp. 108-123, 2014.

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