

Automated Generation of Cross-Domain Analogies via Evolutionary Computation

Atılım Güneş Baydin^{1,2}, Ramon López de Mántaras¹, Santiago Ontañón³

¹Artificial Intelligence Research Institute, IIIA - CSIC

Campus Universitat Autònoma de Barcelona, 08193 Bellaterra, Spain

²Departament d'Enginyeria de la Informació i de les Comunicacions

Universitat Autònoma de Barcelona, 08193 Bellaterra, Spain

³Department of Computer Science, Drexel University, 3141 Chestnut Street, Philadelphia, PA 19104, USA

gunesbaydin@iia.csic.es, mantaras@iia.csic.es, santi@cs.drexel.edu

Abstract

Analogy plays an important role in creativity, and is extensively used in science as well as art. In this paper we introduce a technique for the automated generation of cross-domain analogies based on a novel evolutionary algorithm (EA). Unlike existing work in computational analogy-making restricted to creating analogies between two given cases, our approach, for a given case, is capable of creating an analogy along with the novel analogous case itself. Our algorithm is based on the concept of “memes”, which are units of culture, or knowledge, undergoing variation and selection under a fitness measure, and represents evolving pieces of knowledge as semantic networks. Using a fitness function based on Gentner’s structure mapping theory of analogies, we demonstrate the feasibility of spontaneously generating semantic networks that are analogous to a given base network.

Introduction

In simplest terms, analogy is the transfer of information from a known subject (the *analogue* or *base*) onto another particular subject (the *target*), on the basis of similarity. The cognitive process of analogy is considered at the heart of many defining aspects of human intellectual capacity, including problem solving, perception, memory, and creativity (Holyoak and Thagard 1996); and it has been even argued, by Hofstadter (2001), that analogy is “the core of cognition”.

Analogy-making ability is extensively linked with creative thought (Hofstadter 1995; Holyoak and Thagard 1996; Ward, Smith, and Vaid 2001; Boden 2004) and plays a fundamental role in discoveries and changes of knowledge in arts as well as science, with key examples such as Johannes Kepler’s explanation of the laws of heliocentric planetary motion with an analogy to light radiating from the Sun¹ (Gentner and Markman 1997); or Ernest Rutherford’s analogy between the atom and the Solar System² (Falkenhainer, Forbus, and Gentner 1989). Boden (2004; 2009) classifies analogy as a form of *combinational creativity*, noting that it works by producing unfamiliar combinations of familiar ideas.

In this paper, we present a technique for the automated generation of cross-domain analogies using evolutionary

computation. Existing research on computational analogy is virtually restricted to the discovery and assessment of analogies between a given pair of base case A and target case B (French 2002) (An exception is the Kilaza model by O’Donoghue (2004)). On the other hand, given a base case A, the approach that we present here is capable of creating a novel analogous case B itself, along with the analogical mapping between A and B. This capability of open-ended creation of novel analogous cases is, to our knowledge, the first of its kind and makes our approach highly relevant from a computational creativity perspective. It replicates the psychological observation that an analogy is not always simply “recognized” between an original case and a retrieved analogous case, but the analogous case can sometimes be created together with the analogy (Clement 1988).

As the core of our approach, we introduce a novel evolutionary algorithm (EA) based on the concept of “meme” (Dawkins 1989), where the individuals forming the population represent units of culture, or knowledge, that are undergoing variation, transmission, and selection. We represent individuals as simple semantic networks that are directed graphs of concepts and binary relations (Sowa 1991). These go through variation by memetic versions of EA crossover and mutation, which we adapt to work on semantic networks, utilizing the commonsense knowledge bases of ConceptNet (Havasi, Speer, and Alonso 2007) and WordNet (Fellbaum 1998). Defining a memetic fitness measure using analogical similarity from Gentner’s psychological structure mapping theory (Gentner and Markman 1997), we demonstrate the feasibility of generating semantic networks that are analogous to a given base network.

In this introductory work, we focus on the evolution of analogies using a memetic fitness function promoting analogies. But it is of note that considering different possible fitness measures, the proposed representation and algorithm can serve as a generic tool for the generation of pieces of knowledge with any desired property that is a quantifiable function of the represented knowledge. Our algorithm can also act as a computational model for experimenting with memetic theories of knowledge, such as evolutionary epistemology and cultural selection theory.

After a review of existing research in analogy, evolution, and creativity, the paper introduces details of our algorithm. We then present results and discussion of using the fitness function based on analogical similarity, and conclude with future work and potential applications in creativity.

¹Kepler argued, as light can travel undetectably on its way between the source and destination, and yet illuminate the destination, so can motive force be undetectable on its way from the Sun to planet, yet affect planet’s motion.

²The Rutherford–Bohr model of the atom considers electrons to circle the nucleus in orbits like planets around the Sun, with electrostatic forces providing attraction, rather than gravity.

Background

Analogy

Analogical reasoning has been actively studied from both cognitive and computational perspectives. The dominant school of research in the field, advanced by Gentner (Falkenhainer, Forbus, and Gentner 1989; Gentner and Markman 1997), describes analogy as a structural matching, in which elements from a base domain are mapped to (or aligned with) those in a target domain via structural similarities of their relations. This approach named *structure mapping theory*, with its computational implementation, the *Structure Mapping Engine (SME)* (Falkenhainer, Forbus, and Gentner 1989), has been cited as the most influential work to date on the modeling of analogy-making (French 2002). Alternative approaches in the field include the coherence based view developed by Holyoak and Thagard (Thagard et al. 1990; Holyoak and Thagard 1996), in which analogy is considered as a constraint satisfaction problem involving structure, semantic similarity, and purpose; and the view of Hofstadter (1995) of analogy as a kind of high-level perception, where one situation is perceived as another one. Veale and Keane (1997) extend the work in analogical reasoning to the more specific case of metaphors, which describe the understanding of one kind of thing in terms of another. A highly related cognitive theory is the *conceptual blending* idea developed by Fauconnier and Turner (2002), which involves connecting several existing concepts to create new meaning, operating below the level of consciousness as a fundamental mechanism of cognition. An implementation of this idea is given by Pereira (2007) as a computational model of abstract thought, creativity, and language.

According to whether the base and target cases belong to the same or different domains, there are two types of analogy: *intra-domain*, confined to surface similarities within the same domain; and *cross-domain*, using deep structural similarities between semantically distant information. While much of the research in artificial intelligence has been restricted to intra-domain analogies (e.g. case-based reasoning), studies in psychology have been more concerned with cross-domain analogical experiments (Thagard et al. 1990).

Evolutionary and Memetic Algorithms

Generalizing the mechanisms of the evolutionary process that has given rise to the diversity of life on earth, the approach of *Universal Darwinism* uses a simple progression of variation, natural selection, and heredity to explain a wide variety of phenomena; and it extends the domain of this process to systems outside biology, including economics, psychology, physics, and even culture (Dennett 1995). In terms of application, the metaheuristic optimization method of evolutionary algorithms (EA) provides an implementation of this idea, established as a solid technique with diverse problems in engineering as well as natural and social sciences (Coello Coello, Lamont, and Van Veldhuizen 2007).

In an analogy with the unit of heredity in biological evolution, the gene, the concept of *meme* was introduced by Dawkins (1989) as a unit of idea or information in cultural evolution, hosted, altered, and reproduced in individuals' minds, forming the basis of the field of memetics³.

³Quoting Dawkins (Dawkins 1989): “Examples of memes are tunes, ideas, catch-phrases, clothes fashions, ways of making pots or of building arches. Just as genes propagate themselves in the gene pool by leaping from body to body via sperms or eggs, so memes propagate themselves in the meme pool by leaping from brain to brain...”

Within evolutionary computation, the recently maturing field of memetic algorithms (MA) has experienced increasing interest as a method for solving many hard optimization problems (Moscato, Cotta, and Mendes 2004). The existing formulation of MA is essentially a hybrid approach, combining classical EA with local search, where the population-based global sampling of EA in each generation is followed by an individual learning step mimicking cultural evolution, performed by each candidate solution. For this reason, this approach has been often referred to under different names besides MA, such as “hybrid EA” or “Lamarckian EA”. To date, MA has been successfully applied to a wide variety of problem domains such as NP-hard optimization problems, engineering, machine learning, and robotics.

The potential of an evolutionary approach to creativity has been noted from cultural and practical viewpoints (Gabora 1997; Boden 2009). EA techniques have been shown to emulate creativity in engineering, such as genetic programming (GP) introduced by Koza (2003) as being capable of “routinely producing inventive and creative results”⁴; as well as in visual art, design, and music (Romero and Machado 2008). In psychology, there are studies providing support to an evolutionary view of creativity, such as the behavioral analysis by Simonton (2003) inferring that scientific creativity constitutes a form of constrained stochastic behavior.

The Algorithm

Our approach is based on a *meme pool* comprising individuals represented as semantic networks, subject to variation and selection under a fitness measure. We position our algorithm as a *novel memetic algorithm*, because (1) it is the units of culture, or information, that are undergoing variation, transmission, and selection, very close to the original sense of “memetics” as it was introduced by Dawkins; and (2) this is unlike the existing sense of the word in current MA as an hybridization of individual learning and EA. This algorithm is intended as a new tool focused exclusively on the memetic evolution of knowledge itself, which can find use in knowledge-based systems, reasoning, and creativity.

Our algorithm proceeds similar to a conventional EA cycle (Algorithm 1), with a relatively small set of parameters. We implement semantic networks as linked-list data structures of concept and relation objects. The descriptions of representation, fitness evaluation, variation, and selection steps are presented in the following sections. Parameters affecting each step of the algorithm are given in Table 1.

Algorithm 1 Outline of the algorithm

```
1: procedure MEMETICALGORITHM
2:    $P(t = 0) \leftarrow \text{INITIALIZE}(Pop_{size}, C_{max}, R_{min}, T)$ 
3:   repeat
4:      $\phi(t) \leftarrow \text{EVALUATEFITNESSES}(P(t))$ 
5:      $S(t) \leftarrow \text{SELECTION}(P(t), \phi(t), S_{size}, S_{prob})$ 
6:      $V(t) \leftarrow \text{VARIATION}(S(t), P_c, P_m, T)$ 
7:      $P(t + 1) \leftarrow V(t)$ 
8:      $t \leftarrow t + 1$ 
9:   until stop criterion
10: end procedure
```

⁴Striking examples of demonstrated GP creativity include replication of historically important discoveries in engineering, such as the reinvention of negative feedback circuits originally conceived by Harold Black in 1920s.

Representation

The algorithm is centered on the use of semantic networks (Sowa 1991) for encoding evolving memotypes. An important characteristic of a semantic network is whether it is definitional or assertional: in definitional networks the emphasis is on taxonomic relations (e.g. $IsA(bird, animal)$ ⁵) describing a subsumption hierarchy that is true by definition; in assertional networks, relations describe instantiations that are contingently true (e.g. $AtLocation(human, city)$). In this study we combine the two approaches for increased expressivity. As such, semantic networks provide a simple yet powerful means to represent the “memes” of Dawkins as data structures that are algorithmically manipulatable, allowing a procedural implementation of memetic evolution.

In terms of representation, our approach is similar to several existing graph-based encodings of individuals in EA. The most notable is genetic programming (GP) (Koza et al. 2003), where candidate solutions are computer programs represented in a tree hierarchy. Montes and Wyatt (2004) present a detailed overview of graph-based EA techniques besides GP, which include parallel distributed genetic programming (PDGP), genetic network programming (GNP), evolutionary graph generation, and neural programming.

Using a graph-based representation makes the design of variation operators specific to graphs necessary. In works such as GNP, this is facilitated by using a string-based encoding of node types and connectivity, permitting operators very close to their counterparts in conventional EA; and in PDGP, operations are simplified by making nodes occupy points in a fixed-size two-dimensional grid. What is common with GP related algorithms is that the output of each node in the graph can constitute an input to another node. In comparison, the range of connections that can form a semantic network of a given set of concepts is limited by commonsense knowledge, i.e. the relations have to make sense to be useful (e.g. $IsA(bird, animal)$ is meaningful while $Causes(bird, table)$ is not). To address this issue, we introduce new crossover and mutation operations for memetic variation, making use of commonsense reasoning (Mueller 2006) and adapted to work on semantic networks.

Commonsense Knowledge Bases Commonsense reasoning refers to the type of reasoning involved in everyday thinking, based on *commonsense knowledge* that an ordinary person is expected to know, or “the knowledge of how the world works” (Mueller 2006). Knowledge bases such as the *ConceptNet*⁶ project of MIT Media Lab (Havasi, Speer, and Alonso 2007) and *Cyc*⁷ maintained by Cycorp company are set up to assemble and classify commonsense information. The lexical database *WordNet*⁸ maintained by the Cognitive Science Laboratory at Princeton University also has characteristics of a commonsense knowledge base, via synonym, hypernym⁹, and hyponym¹⁰ relations (Fellbaum 1998).

In our implementation we make use of ConceptNet version 4 and WordNet version 3 to process commonsense knowledge, where ConceptNet contributes around 560,000

definitional and assertional relations involving 320,000 concepts and WordNet contributes definitional relations involving around 117,000 synsets¹¹. The hypernym and hyponym relations among noun synsets in WordNet provide a reliable collection of *IsA* relations. In contrast, the variety of assertions in ConceptNet, contributed by volunteers across the world, makes it more prone to noise. We address this by ignoring all assertions with a reliability score (determined by contributors’ voting) below a set minimum R_{min} (Table 1).

Initialization

At the start of each run of the algorithm, the population of size Pop_{size} is initialized with individuals created by *random semantic network generation* (Algorithm 1). This is achieved by starting from a network comprising only one concept randomly picked from commonsense knowledge bases and running a semantic network expansion algorithm that (1) randomly picks a concept in the given network (e.g. *human*); (2) compiles a list of relations—from commonsense knowledge bases—that the picked concept can be involved in (e.g. $\{CapableOf(human, think), Desires(human, eat), \dots\}$); (3) appends to the network a relation randomly picked from this list, together with the other involved concept; and (4) repeats this until a given number of concepts has been appended or a set timeout T has been reached (covering situations where there are not enough relations). Note that even if grown in a random manner, the resulting network itself is totally meaningful and consistent because it is a combination of rational information from commonsense knowledge bases.

The initialization algorithm depends upon the parameters of C_{max} , the maximum number of initial concepts, and R_{min} , the minimum ConceptNet relation score (Table 1).

Fitness Measure

Since the individuals in our approach represent knowledge, or memes, the fitness for evolutionary selection is defined as a function of the represented knowledge. For the automated generation of analogies through evolution, we introduce a memetic fitness based on analogical similarity with a given semantic network, utilizing the Structure Mapping Engine (SME) (Falkenhainer, Forbus, and Gentner 1989; Gentner and Markman 1997). Taking the analogical matching score from SME as the fitness, our algorithm can evolve collections of information that are analogous to a given one.

In SME, an analogy is a one-to-one mapping from the base domain into the target domain, which correspond, in our fitness measure, to the semantic network supplied at the start and the individual networks whose fitnesses are evaluated by the function. The mapping is guided by the structure of relations between concepts in the two domains, ignoring the semantics of the concepts themselves; and is based on the systematicity principle, where connected knowledge is preferred over independent facts and is assigned a higher matching score. As an example, the Rutherford–Bohr atom and Solar System analogy (Gentner and Markman 1997) would involve a mapping from *sun* and *planet* in the first domain to *nucleus* and *electron* in the second domain. The labels and structure of relations in the two domains (e.g. $\{Attracts(sun, planet), Orbits(planet, sun), \dots\}$ and $\{Attracts(nucleus, electron), Orbits(electron,$

⁵Here we adopt the notation $IsA(bird, animal)$ to mean that the concepts *bird* and *animal* are connected by the directed relation *IsA*, i.e. “bird is an animal.”

⁶<http://conceptnet.media.mit.edu>

⁷<http://www.cyc.com>

⁸<http://wordnet.princeton.edu>

⁹*Y* is a *hypernym* of *X* if every *X* is a (kind of) *Y* ($IsA(dog, canine)$).

¹⁰*Y* is a *hyponym* of *X* if every *Y* is a (kind of) *X*.

¹¹A *synset* is a set of synonyms that are interchangeable without changing the truth value of any propositions in which they are embedded.

nucleus), \dots) define and constrain the possible mappings between concepts that can be formed by SME.

We make use of our own implementation of SME based on the original description by Falkenhainer et al. (1989) and adapt it to the simple concept–relation structure of semantic networks, by mapping the predicate calculus constructs of *entities* into *concepts*, *relations* to *relations*, *attributes* to *ISA relations*, and excluding *functions*.

Selection

After assigning fitness values to all individuals in the current generation, these are replaced with offspring generated by variation operators on parents. The parents are probabilistically selected from the population according to their fitness, with reselection allowed. While individuals with a higher fitness have a better chance of being selected, even those with low fitness have a chance to produce offspring, however small. In our experiments we employ tournament selection (Coello Coello, Lamont, and Van Veldhuizen 2007), meaning that for each selection, a “tournament” is held among a few randomly chosen individuals, and the more fit individual of each successive pair is the winner according to a winning probability (Table 1).

In each cycle of algorithm, crossover is applied to parents selected from the population until $Pop_{size} \times P_c$ offspring are created (Table 1). Mutation is applied to $Pop_{size} \times P_m$ selected individuals, supplying the remaining part of the next generation (i.e. $P_c + P_m = 1$). We also employ elitism, by replacing a randomly picked offspring in next generation with the individual with the current best fitness.

Variation Operators

In contrast with existing graph-based evolutionary approaches that we have mentioned, our representation does not permit arbitrary connections between different nodes and requires variation operators that should be based on information provided by commonsense knowledge bases. This means that any variation operation on the individuals should: (1) preserve the structure within boundaries set by commonsense knowledge; and (2) ensure that even vertices and edges randomly introduced into a semantic network connect to existing ones through meaningful and consistent relations¹².

Here we present commonsense crossover and mutation operators specific to semantic networks.

Commonsense Crossover In classical EA, features representing individuals are commonly encoded as linear strings and the crossover operation simulating genetic recombination is simply a cutting and merging of this one dimensional object from two parents. In graph-based approaches such as GP, subgraphs can be freely exchanged between parent graphs (Koza et al. 2003; Montes and Wyatt 2004). Here, as mentioned, the requirement that a semantic network has to make sense imposes significant constraints on the nature of recombination.

We introduce two types of *commonsense crossover* that are tried in sequence by the variation algorithm. The first type attempts a sub-graph interchange between two selected parents similar to common crossover in standard GP; and where this is not feasible due to the commonsense structure of relations forming the parents, the second type falls back to a combination of both parents into a new offspring.

¹²It should be noted that we depend on the meaningfulness and consistency (i.e. compatibility of relations with others involving the same concepts) of information in the commonsense knowledge bases, which should be ensured during their maintenance.

Type I (subgraph crossover): A pair of concepts, one from each parent, that are *interchangeable*¹³ are selected as *crossover concepts*, picked randomly out of all possible such pairs. For instance, in Figure 1, *bird* and *airplane* are interchangeable, since they can replace each other in the relations *CapableOf*(\cdot , *fly*) and *AtLocation*(\cdot , *air*). In each parent, a subgraph is formed, containing: (1) the crossover concept; (2) the set of all relations, and associated concepts, that are not common with the other crossover concept (In Figure 1 (a), *HasA*(*bird*, *feather*) and *AtLocation*(*bird*, *forest*); and in (b) *HasA*(*airplane*, *propeller*), *MadeOf*(*airplane*, *metal*), and *UsedFor*(*airplane*, *travel*)); and (3) the set of all relations and concepts connected to these (In Figure 1 (a) *PartOf*(*feather*, *wing*) and *PartOf*(*tree*, *forest*); and in (b) *MadeOf*(*propeller*, *metal*)), excluding the ones that are also one of those common with the other crossover concept (the concept *fly* in Figure 1 (a), because of the relation *CapableOf*(\cdot , *fly*)). This, in effect, forms a subgraph of information specific to the crossover concept, which is insertable into the other parent. Any relations between the subgraph and the rest of the network not going through the crossover concept are severed (e.g. *UsedFor*(*wing*, *fly*) in Figure 1 (a)). The two offspring are formed by exchanging these subgraphs between the parent networks (Figure 1 (c) and (d)).

Type II (graph merging crossover): A concept from each parent that is *attachable*¹⁴ to the other parent is selected as a *crossover concept*. The two parents are merged into an offspring by attaching a concept in one parent to another concept in the other parent, picked randomly out of all possible attachments (*CreatedBy*(*art*, *human*) in Figure 2. Another possibility is *Desires*(*human*, *joy*)). The second offspring is formed randomly the same way. In the case that no attachable concepts are found, the parents are merged as two separate clusters within the same semantic network.

Commonsense Mutation We introduce several types of *commonsense mutation* operators that modify a parent by means of information from commonsense knowledge bases. For each mutation to be performed, the type is picked at random with uniform probability. If the selected type of mutation is not feasible due to the commonsense structure of the parent, another type is again picked. In the case that a set timeout of T trials has been reached without any operation, the parent is returned as it is.

Type I (concept attachment): A new concept randomly picked from the set of concepts *attachable* to the parent is attached through a new relation to one of existing concepts (Figure 3 (a) and (b)).

Type IIa (relation addition): A new relation connecting two existing concepts in the parent is added, possibly connecting unconnected clusters within the same network (Figure 3 (c) and (d)).

Type IIb (relation deletion): A randomly picked relation in the parent is deleted, possibly leaving unconnected clusters within the same network (Figure 3 (e) and (f)).

Type IIIa (concept addition): A randomly picked new

¹³We define two concepts from different semantic networks as *interchangeable* if both can replace the other in all, or part, of the relations the other is involved in, queried from commonsense knowledge bases.

¹⁴We define a distinct concept as *attachable* to a semantic network if at least one commonsense relation connecting the concept to any of the concepts in the network can be discovered from commonsense knowledge bases.

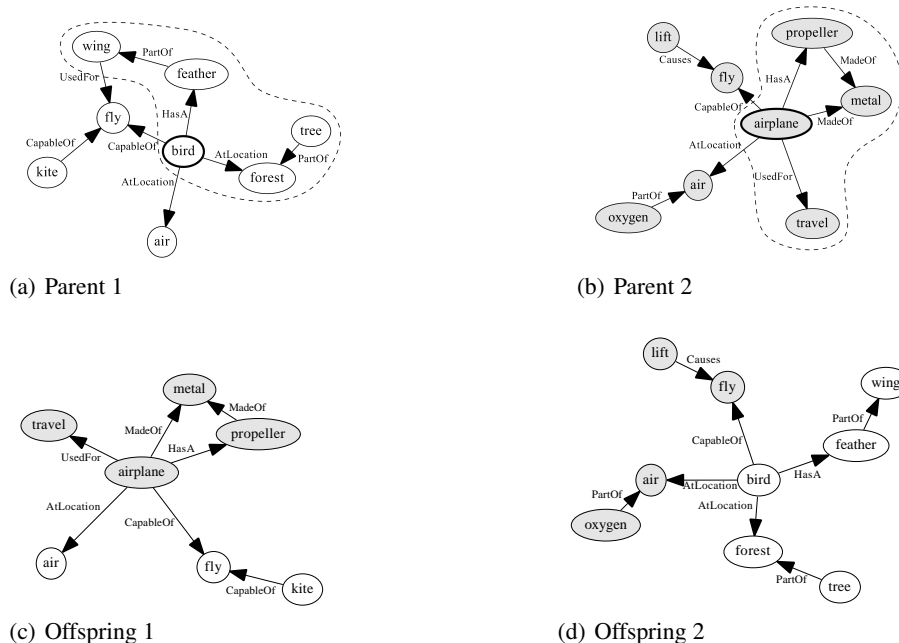


Figure 1: Commonsense crossover type I (subgraph crossover), centered on the concepts of *bird* for parent 1 and *airplane* for parent 2.

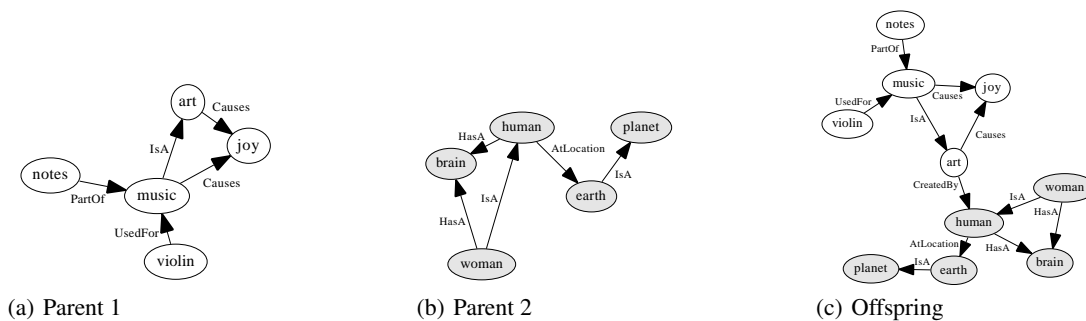


Figure 2: Commonsense crossover type II (graph merging crossover), merging by the relation $CreatedBy(art, human)$. If no concepts attachable through commonsense relations are encountered, the offspring is formed by merging the parent networks as two separate clusters within the same semantic network.

concept is added to the parent as a new cluster (Figure 3 (g) and (h)).

Type IIIb (concept deletion): A randomly picked concept is deleted with all the relations it is involved in, possibly leaving unconnected clusters within the same network (Figure 3 (i) and (j)).

Type IV (concept replacement): A concept in the parent, randomly picked from the set of those with at least one *interchangeable* concept, is replaced with one (randomly picked) of its interchangeable concepts. Any relations left unsatisfied by the new concept are deleted (Figure 3 (k) and (l)).

Results and Discussion

In this introductory study, we adopt values for crossover and mutation probabilities similar to earlier studies in graph-based EA (Koza et al. 2003; Montes and Wyatt 2004) (Ta-

ble 1). We use a crossover probability of $P_c = 0.85$, and a somewhat-above-average mutation rate of $P_m = 0.15$, accounting for the high tendency of mutation postulated in memetic literature¹⁵. In our experiments, we subject a population of $Pop_{size} = 200$ individuals to tournament selection with tournament size $S_{size} = 8$ and winning probability $S_{prob} = 0.8$.

Using this parameter set, we present the results from two runs of experiment: evolved analogies for a network describing some basic astronomical knowledge are shown in Figure 4 and for a network of familial relations in Figure 5. We show in Figure 6 (a) the progress of the best and average fitness in the population during the run that produced the results in Figure 4. The best and average size of semantic

¹⁵See Gil-White (Gil-White 2008) for a review and discussion of mutation in memetics.




Figure 3: Examples illustrating the types of commonsense mutation used in this study.

networks forming the individuals are shown in Figure 6 (b). We observe that evolution asymptotically reaches a fitness plateau after about 40 generations. This coincides roughly with the point where the size of the best individual (13–14) becomes comparable with that of the given base semantic network (11, in Figure 4), after which improvements in the one-to-one analogy become sparser and less feasible. We also note that, between generations 21–34, the best network size actually gets smaller, demonstrating the possibility of improvement in network configuration without adding further nodes. Our experiments demonstrate that the proposed algorithm is capable of spontaneously creating collections of knowledge analogous to the one given in a base semantic network, with very good performance. In most cases, our implementation was able to reach extensive analogies within 50 generations and reasonable computational time.

From an analogical reasoning viewpoint, the algorithm achieves the generation of diverse novel cases analogous to a given case. Compared with the Kilaza model of O’Donoghue (2004) for finding novel analogous cases, which works by evaluating possible analogies to a given target case from a collection of candidate source domains that are assumed to be available, our approach is capable of open-ended and spontaneous creation of analogous cases from the ground up, replicating an essential mode of creative behavior observed in psychology (Clement 1988).

An important result is that, even if the use of commonsense knowledge in our algorithm was prompted by concerns that are practical in nature (i.e. restrictions on the meaningfulness and consistency of memetic variation by the introduced crossover and mutation operators), it eventually serves to tackle a very fundamental and long-standing

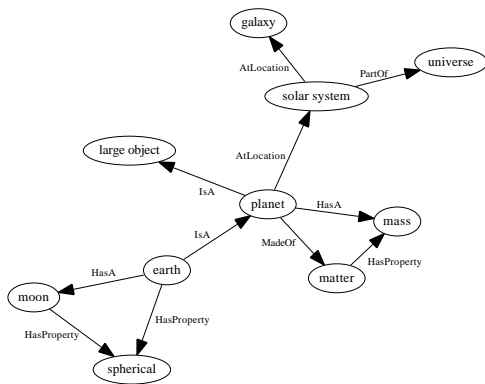
Table 1: Parameters used during experiments

	Parameter	Value
Evolution	Population size (P_{pop_size})	200
	Crossover probability (P_c)	0.85
	Mutation probability (P_m)	0.15
Semantic networks	Max. initial concepts (C_{max})	5
	Min. relation score (R_{min})	2.0
	Timeout (T)	10
Selection	Type	Tournament
	Tournament size (S_{size})	8
	Tournament win prob. (S_{prob})	0.8
	Elitism	Employed

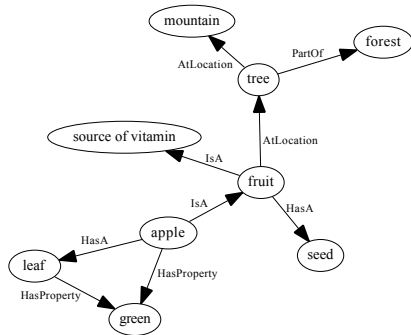
problem in computational creativity: as put forth by Boden (2009), “no current AI system has access to the rich and subtly structured stock of concepts that any normal adult human being has built up over a lifetime” and “what’s missing, as compared with the human mind, is the rich store of world knowledge (including cultural knowledge) that’s often involved.” We believe that the inherent commonsense reasoning element in our approach provides a means to address this criticism of lack of world knowledge in computational approaches to creativity.

Conclusions and Future Work

We have presented a novel evolutionary algorithm that employs semantic networks as evolving individuals, paralleling



(a) Given semantic network, 10 concepts, 11 relations (base domain)




(b) Evolved individual, 9 concepts, 9 relations (target domain)

Figure 4: Experiment 1: The evolved individual is encountered after 35 generations, with fitness value 2.8. Concepts and relations of the individual not involved in the analogy are not shown here for clarity.


the model of cultural evolution in the field of memetics. This algorithm, to our knowledge, is the first of its kind. The use of semantic networks provides a suitable basis for implementing variation and selection of memes as put forth by Dawkins (Dawkins 1989). We have introduced preliminary versions of variation operators that work on this representation, utilizing knowledge from commonsense knowledge bases. We have also contributed a memetic fitness measure based on the structure mapping theory from psychology.

Even if it is an intuitive fact that human culture and knowledge are evolving with time, existing models of culture, in their current state, are too minimalistic and weak in their descriptions of individual creativity and novelty; and conversely, theories modeling individual creativity lack consideration of cultural transmission and replication (Gabora 1997). We believe that studies exploring creativity with evolutionary approaches have the potential for bridging this gap.

In future work, an interesting possibility is to start the random semantic network generation procedure with several given concepts, allowing the discovery of cases formed around a particular set of seed concepts. The simple fitness function used in this introductory study can be extended to take graph-theoretical properties of semantic networks into account, such as the number of nodes or edges, shortest path length, or the clustering coefficient. The research would also benefit from exploring different types of mutation and crossover, and grounding the design of such operators on



(a) Given semantic network, 11 concepts, 11 relations (base domain)



(b) Evolved individual, 10 concepts, 9 relations (target domain)

Figure 5: Experiment 2: The evolved individual is encountered after 42 generations, with fitness value 2.7. Concepts and relations of the individual not involved in the analogy are not shown here for clarity.

existing theories of cultural transmission and variation, discussed in sociological theories of knowledge.


A direct and very interesting application of our approach would be to devise experiments with realistically formed fitness functions modeling selectionist theories of knowledge, which remain untested until this time. One such theory is the *evolutionary epistemology* of Campbell (Bickhard and Campbell 2003), describing the development of human knowledge and creativity through selectionist principles such as blind variation and selective retention (BVSr).

Acknowledgments


This work was supported by a JAE-Predoc fellowship from CSIC, and the research grants: 2009-SGR-1434 from the Generalitat de Catalunya, CSD2007-0022 from MICINN, and Next-CBR TIN2009-13692-C03-01 from MICINN.

References

- [Bickhard and Campbell 2003] Bickhard, M. H., and Campbell, D. T. 2003. Variations in variation and selection: the ubiquity of the variation-and-selective-retention ratchet in emergent organizational complexity. *Foundations of Science* 8:215–2182.
- [Boden 2004] Boden, M. A. 2004. *The Creative Mind: Myths and Mechanisms*. London: Routledge, second edition.



(a)



(b)

Figure 6: Evolution of (a) fitness and (b) semantic network size during the course of an experiment with parameters given in Table 1. Filled circles represent the best individual in a generation, empty circles represent population average. Network size is taken to be the number of relations (edges).

- [Boden 2009] Boden, M. A. 2009. Computer models of creativity. *AI Magazine* 30(3):23–34.
- [Clement 1988] Clement, J. 1988. Observed methods for generating analogies in scientific problem solving. *Cognitive Science* 12:563–586.
- [Coello Coello, Lamont, and Van Veldhuizen 2007] Coello Coello, C. A.; Lamont, G. B.; and Van Veldhuizen, D. A. 2007. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Springer.
- [Dawkins 1989] Dawkins, R. 1989. *The Selfish Gene*. Oxford University Press.
- [Dennett 1995] Dennett, D. C. 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. Simon & Schuster.
- [Falkenhainer, Forbus, and Gentner 1989] Falkenhainer, B.; Forbus, K. D.; and Gentner, D. 1989. The Structure-Mapping Engine: Algorithm and examples. *Artificial Intelligence* 41:1–63.
- [Fauconnier and Mark 2002] Fauconnier, G., and Mark, T. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.
- [Fellbaum 1998] Fellbaum, C. 1998. *WordNet: An Electronic Lexical Database*. MIT Press.
- [French 2002] French, R. M. 2002. The computational modeling of analogy-making. *Trends in Cognitive Sciences* 6(5):200–205.
- [Gabora 1997] Gabora, L. 1997. The origin and evolution of

culture and creativity. *Journal of Memetics – Evolutionary Models of Transmission* 1.

- [Gentner and Markman 1997] Gentner, D., and Markman, A. B. 1997. Structure mapping in analogy and similarity. *American Psychologist* 52:45–56.
- [Gil-White 2008] Gil-White, F. 2008. Let the meme be (a meme): insisting too much on the genetic analogy will turn it into a straightjacket. In Botz-Bornstein, T., ed., *Culture, Nature, Memes*. Newcastle upon Tyne: Cambridge Scholars.
- [Havasi, Speer, and Alonso 2007] Havasi, C.; Speer, R.; and Alonso, J. 2007. ConceptNet 3: a flexible, multilingual semantic network for common sense knowledge. In *Proceedings of Recent Advances in Natural Language Processing*.
- [Hofstadter 1995] Hofstadter, D. R. 1995. *Fluid concepts and creative analogies: Computer models of the fundamental mechanisms of thought*. New York: Basic Books.
- [Hofstadter 2001] Hofstadter, D. 2001. Analogy as the core of cognition. In Gentner, D.; Holyoak, K. J.; and Kokinov, B., eds., *Analogical Mind: Perspectives From Cognitive Science*. Cambridge, MA: MIT Press. 499–538.
- [Holyoak and Thagard 1996] Holyoak, K. J., and Thagard, P. 1996. *Mental Leaps: Analogy in Creative Thought*. Bradford Books.
- [Koza et al. 2003] Koza, J. R.; Keane, M. A.; Streeter, M. J.; Mydlowec, W.; Yu, J.; and Lanza, G. 2003. *Genetic Programming IV: Routine Human-Competitive Machine Intelligence*. Kluwer Academic Publishers.
- [Montes and Wyatt 2004] Montes, H. A., and Wyatt, J. L. 2004. Graph representation for program evolution: An overview. Technical report, University of Birmingham School of Computer Science.
- [Moscato, Cotta, and Mendes 2004] Moscato, P. A.; Cotta, C.; and Mendes, A. 2004. *Studies in Fuzziness and Soft Computing – New Optimization Techniques in Engineering*. New York: Springer. chapter Memetic Algorithms.
- [Mueller 2006] Mueller, E. T. 2006. *Commonsense Reasoning*. Morgan Kaufmann.
- [O'Donoghue 2004] O'Donoghue, D. 2004. *Finding Novel Analogies*. Ph.D. Dissertation, University College Dublin, Department of Computer Science, Dublin, Ireland.
- [Pereira 2007] Pereira, F. C. 2007. *Creativity and Artificial Intelligence: A Conceptual Blending Approach*. New York: Mouton de Gruyter.
- [Romero and Machado 2008] Romero, J. J., and Machado, P. 2008. *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*. Springer.
- [Simonton 2003] Simonton, D. K. 2003. Scientific creativity as constrained stochastic behavior: The integration of product, person, and process perspectives. *Psychological Bulletin* 129(4):475–494.
- [Sowa 1991] Sowa, J. F. 1991. *Principles of Semantic Networks: Explorations in the Representation of Knowledge*. San Mateo: Morgan Kaufmann.
- [Thagard et al. 1990] Thagard, P.; Holyoak, K. J.; Nelson, G.; and Gochfeld, D. 1990. Analog retrieval by constraint satisfaction. *Artificial Intelligence* 46:259–310.
- [Veale and Keane 1997] Veale, T., and Keane, M. 1997. The competence of sub-optimal structure mapping on hard analogies. In *Proceedings of the 15th International Joint Conference on AI*. San Mateo, CA: Morgan Kaufmann.
- [Ward, Smith, and Vaid 2001] Ward, T. B.; Smith, S. M.; and Vaid, J. 2001. *Creative Thought: An Investigation of Conceptual Structures and Processes*. Washington, DC: American Psychological Association.