

Going Into Greater Depth in the Quest for Hidden Frames

João Gonçalves^{*1}, Aaron Bembek², Pedro Martins¹, and Amílcar Cardoso¹

¹CISUC, Department of Informatics Engineering, University of Coimbra, Portugal

²Harvard University, Cambridge, Massachusetts, USA

{jgconc@dei.uc.pt, bembek@g.harvard.edu, pjmm@dei.uc.pt, amilcar@dei.uc.pt}

Abstract

Semantic frames are a fundamental ingredient in computational implementations of Conceptual Blending (CB) theory. Therefore, we may ask the question of how to build them or where to retrieve them. This paper offers a solution which is to explore large-semantic networks for repeating structures resembling frames. We also include a feature the frames could have to give them a sense of completeness. Potential patterns were searched with a Multi-Objective-Evolutionary-Algorithm (MOEA) giving wider and ampler results when compared with a Single-Objective stochastic search.

Introduction

The Conceptual Blending (CB) framework (Fauconnier and Turner 2002) was suggested as a cognitive theory interpreting how conceptual integration processes occur in the human mind, as well as the creation of meaning, argumentation and the transmission of ideas (Coulson 2006). Although not devised by its authors to explain the formation of creative constructs, CB has been successfully used as the main engine in many Computational Creativity (CC) systems such as (Pereira 2005; Gonçalves, Martins, and Cardoso 2017; Cunha et al. 2017; Eppe et al. 2018; Martins et al. 2019).

The theory involves interactions between four mental spaces: two *input spaces*, a *generic space* and the *blend space* (Fig. 1). The latter contains the “output” of the CB process. Each mental space corresponds to a partial and temporary structure of knowledge built for the purpose of local understanding and action (Fauconnier 1994). In some implementations of CB, including ours, the mental spaces are stored as semantic graphs. These are networks of vertices (the concepts) connected by directed edges (the relations). Each relation/edge states a fact between a subject and an object such as:

partOf(wing, bird).

CB theory also mentions frames which are required in some computational models of CB (Pereira 2005), including one we have been working on (Gonçalves, Martins, and Cardoso 2017). They are needed to guide the blending process towards stable and recognisable wholes. Frames represent situations or interactions involving various participants and

*Corresponding author.

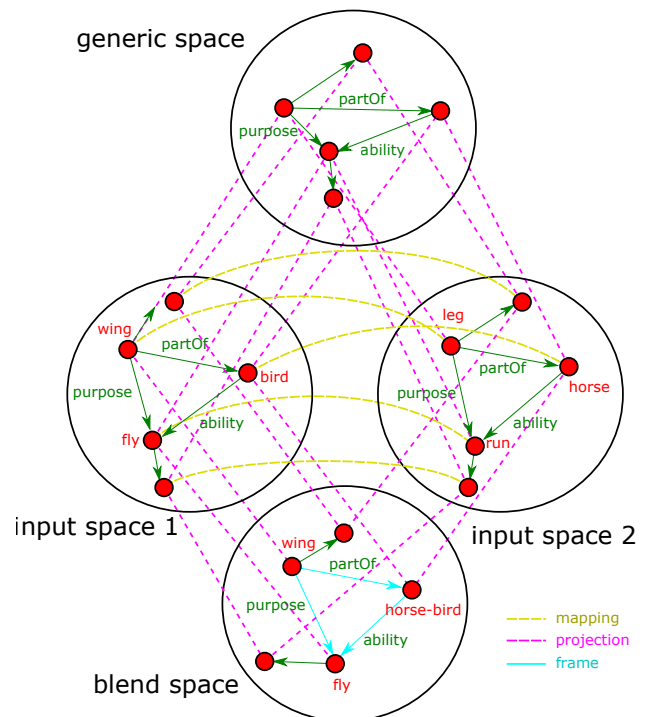


Figure 1: The four mental spaces of CB theory with two examples of input spaces: horse and bird. One possible blend of these input spaces could slightly resemble the *Pegasus*.

can be very general as well as very specific (Johnson and Fillmore 2000). Frames can be organised either as a lattice or a taxonomic structure. For example, within the domain of *motion*, the *transportation* frame provides *movers* with *means* of transportation along a *path* (Baker, Fillmore, and Lowe 1998). A more intuitive example is the frame marriage where, amongst other properties, two people have witnesses and share vows. In this case, a possible mental space containing this frame would be Mary’s marriage. In practical terms, a semantic frame is composed of either specific or abstract concepts connected by relations between them with the whole representing a meaningful entity, event or other abstract composite concept.

We now state the motivation for the current work. The first is the intention of finding useful and interesting frames on large-scale networks. The subject of finding appealing semantic frames was previously addressed with a computational model (Gonçalves, Martins, and Cardoso 2018), being the system described in this paper a direct descendant. Additionally, we were wondering if it was possible to find structures resembling frames in the same Knowledge Base (KB) functioning as the source of input spaces. Secondly, although some repositories of frames do exist - FrameNet (Ruppenhofer et al. 2016), MetaNet (David et al. 2014) and Framester (Gangemi et al. 2016), amongst others - we do think they are better aimed at linguistic systems and not easily usable by computational models of CB. Hence our second motivation, the building of a sufficiently comprehensive repository of frames to help with computational implementations of CB. The third motivation is the implementation of visual tools to help researchers who work with semantic frames in general.

After this introductory section, we follow with a short description of the CB theory, then with the importance of semantic frames in CB and our latest approach to frame mining. Later, we present and discuss the results and conclude on our findings, followed by what we expect as further work.

Conceptual Blending (CB) Theory

The input spaces serve as the initial sources of knowledge and supply the content that will be blended. A partial mapping is first established between both input spaces, reflecting a sense of similarity between them. This mapping associates the input spaces and is mirrored in the generic space, encapsulating the elements shared by the input spaces. A custom selection of this mapping is used to partially project structures from both input spaces - including nearby elements - integrating them in an emerging structure called the blend.

The integration of input elements from the input spaces (Fauconnier and Turner 1998) in the blend space is split in three sub-tasks: **composition**, **completion** and **elaboration**. The first is the projection of elements from the input spaces into the blend space. Completion corresponds to the use of existing knowledge in the form of background frames and the generation of meaningful structures in the blend. Elaboration performs cognitive work in the blend according to its ongoing emergent structures. The order of these tasks does not need to be pre-determined and several iterations may be necessary (Pereira and Cardoso 2003).

The CB process allows for substantial diversity of generated blends that - depending on the quantity of knowledge being handled - may be computationally unbounded (Martins et al. 2016). To guide the integration process towards highly integrated, coherent and easily interpreted blends, (Fauconnier and Turner 2002) proposed eight **optimality principles**. We outline two principles stating the relevance of frames according to Fauconnier and Turner:

Integration - the blend must be perceived and manipulated as a unit. Every element in the blend should have integration.

Pattern Completion - elements in the blend should be completed by using existing patterns as additional inputs. A

completing frame should be used that has compressed versions of important vital relations between the inputs.

Hunting for Semantic Frames

The frames are assumed to be patterns composed of more than one relation between three or more concepts. An example is seen in Fig. 1 as the three connected relations *purpose*, *partOf* and *ability*. The system searches for recurring patterns such as those in a KB that contains semantic graphs representing the input spaces. The concepts present in the patterns are converted to variables (words starting with a Capital letter) and a Prolog query is made from the pattern. Using the example shown in Fig. 1 this would be the following query:

ability(A, C), purpose(B, C), partOf(B, A).

A frame is satisfied if instantiating its variables with different concepts present in the KB, the frame's conditions agree with the KB's facts. The number of unique possibilities for these variables (as well as their combination) represents the prevalence of the frame's structure in the KB. We see this as an important factor and it is one of multiple objectives to be solved. The remaining objectives are explained in the following section.

Multi Objective Evolutionary Algorithm (MOEA)

The search for patterns resembling frames is handled by a MOEA evolving a set of chromosomes where each encodes a pattern. These are mutated according to the relations existing in the KB, adequately adding relations from the KB or removing existing ones from the pattern. The mutation is done while maintaining consistency between the relations in the pattern and in the KB.

In (Gonçalves, Martins, and Cardoso 2018), the fitness of each pattern was a weighted sum of various objectives. It is now a set of various competing objectives. This allows for a further exploration of the search space, covering a higher diversity of frames within the range of all the objectives. A GUI (Fig. 2) was also implemented to help with the visualisation of the Non-Dominated Set (the solutions not dominated by others).

Three objectives were outlined, all to be maximised. The first is the number of solutions a frame has. The second is the number of unique labels existing in the pattern's relations. The third is an idea we named *ucycles* (short for undirected cycles). The stochastic algorithm of the previous work rarely found patterns with *ucycles* (at most one in a million patterns) and thus the justification for an additional objective. Our reason for frames with *ucycles* is more of a subjective one, but we think that those frames have a sense of completeness, because the concepts in a *ucycle* are inter-related and closed as a whole.

A *ucycle* of a graph is a path of *distinct* edges where any vertex is reachable from itself, ignoring the direction of the edges. The number of *ucycles* a pattern contains is calculated using an adapted depth-first expansion. It goes through all the connected vertices of the pattern while skipping previously expanded relations to prevent the algorithm from

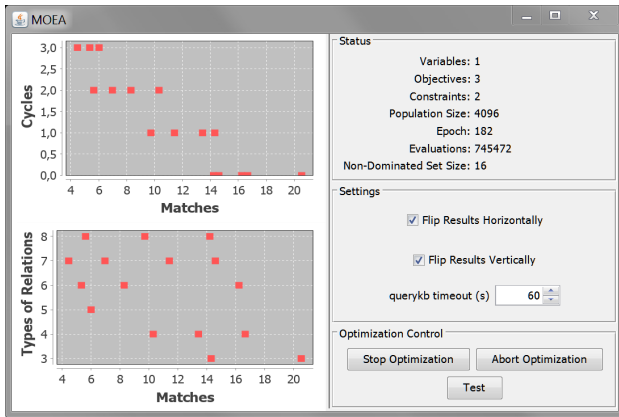


Figure 2: Graphical User Interface of the MOEA with three objectives and the Non-Dominated Set (red squares).




Figure 3: Example of a frame with an undirected cycle.

getting stuck. The expansion traverses all connected edges existing in the pattern, regardless of their direction (Fig. 3).

Counting the Materialising Frames

Our multithreaded *querykb* tool counts the number of solutions satisfying a possible frame, one of the objectives in our MOEA. The number of occurrences of a pattern can be immense, and many patterns occur too many times to be counted using a 64-bit integer. Instead, the count is internally represented as a Big Integer. The logarithm to the base 10 of this Big Integer is then returned to the MOEA. Since we only care about the number of solutions and not what they actually are, we can often avoid explicitly enumerating all of them. Consider the simple conjunctive query

$$p(U, V), q(X, Y).$$

A naive approach to counting the number of solutions to this query would explicitly enumerate all substitutions ranging over $U, V, X,$ and Y that make this conjunction true. However, to count the number of solutions we can just multiply the cardinality of p by the cardinality of q . Our tool generalises this intuition to handle cases where variables are shared between conjuncts. As it evaluates a conjunctive query, it maintains an intermediate relation that pairs substitutions with integers. The integer associated with a substitution represents how many substitutions have been merged into that single substitution. Substitutions are merged when they are equivalent modulo bindings for variables that do not appear later in the query. For example, when the tool evalu-

ates the first conjunct of the query

$$p(X, Y), q(Y, Z).$$

it will produce an intermediate relation with n entries, where n is the number of distinct values of Y such that $p(X, Y)$ holds for some X . Each entry associates a substitution for Y with the number of bindings for X such that $p(X, Y)$ holds for that fixed value of Y ; that is, the entry represents the result of merging together all the discovered substitutions ranging over X and Y that have that binding for Y .

To maximise the opportunities to merge substitutions, we use a query planning heuristic that exploits the structure of the graph representation of the query. In this undirected graph, vertices represent variables and edges signify that two variables appear in the same conjunct. Our heuristic searches for a bridge E whose endpoints have high degrees (a bridge is an edge whose deletion leads to a larger number of components). Call the component at one end of the bridge A and the component at the other end of the bridge B . The heuristic constrains the query evaluation plan so that all the conjuncts that appear in A are evaluated before the conjunct represented by E , and all the conjuncts that appear in B are evaluated after the conjunct represented by E . This means that by the time query evaluation reaches the conjunct represented by E , the bindings for all the variables in A (except the endpoint of E) have been merged away, as these variables cannot appear in any conjunct in B . The same heuristic is then recursively applied to the components A and B . For a component containing no bridges, we order the conjuncts using a heuristic that eagerly minimizes the domain of substitutions in intermediate relations.

Results and Discussion

The KB supplying the facts was a custom version of ConceptNet V5 (Speer and Havasi 2012) with 1 229 508 concepts and 1 791 604 relations. We removed from the KB four relations: *isa*, *derivedfrom*, *synonym* and *similar*to. In our opinion, these relations are very generic and do not seem to be fruitful for the CB process. Without these the KB had 35 types of relations.

We used the MOEA Framework (Hadka 2015) with NSGA-II (Deb et al. 2000) as the evolutionary algorithm. The population size was 4096 chromosomes per epoch. The MOEA was executed on a machine with two Intel Xeon eight-core E5-2667v2 processors and 64 GB of RAM. JVM's heapsize was set to a maximum of 48 GB. The *querykb* tool used a block size of 256, 32 threads and a processing time limit for each pattern of five minutes. Five experiments were executed averaging 48 ± 24 hours and 700 ± 200 epochs per experiment.

We accumulated all the experiments in single dataset with 90 964 patterns. Given the colossal amount of patterns we developed a graphical tool to help with both the filtering and selection of promising frames (Fig. 4). We did not find on the web a graphical tool allowing a visualisation of such a large number of graphs, which inspired us to create ours. It shows the patterns as semantic graphs, allows their sorting according to both the objectives and properties of graphs and filtering patterns within a given range of properties or

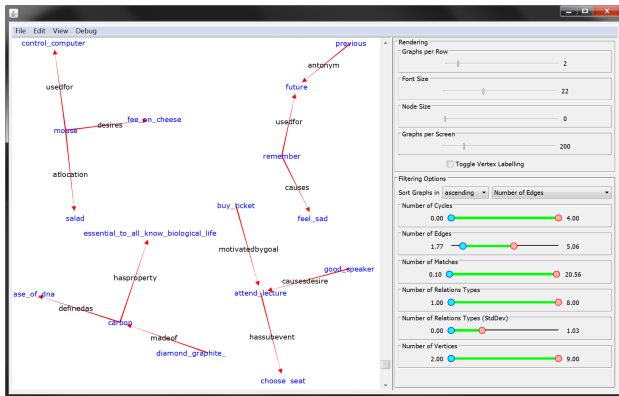


Figure 4: The semantic graph visualisation tool which also serves as a filtering aid for frames.

objectives. The tool supports the rendering of most types of semantic graphs and will be made freely available¹.

We now disclose a few patterns that we believe are interesting, including humorous ones. The patterns' variables are instantiated with examples of concepts (shown in blue) which fully satisfy the pattern in order to be easier to understand. The patterns can be seen in Figs. 5 and 6 with both figures differing on the existence of *ucycles*. The composition of the patterns is highly variable, mainly regarding the relations but we think that they can be used as representation of frames, given their recognisable structure. All these patterns had at least 1 000 occurrences in the KB.




Figure 5: Example of three patterns from the experiments.




Figure 6: Three patterns with *ucycles*, the first two with one *ucycle* and the rightmost pattern with two.

The pattern in Fig. 5a may represent a frame where a place (*mesopotamia*) has a given spoken language (*akkadian language*), partitioned in sub-regions (one of which is *babylon*) with some element (the goddess *Ningal*) belonging to

¹<https://github.com/jcgonc>

the region's context. Fig. 5b shows a pattern depicting an entity (*sulfur*) made of entity (*sulfide*) being part of a larger object (*coal*) found at a specific place (*chemistry lab*). Finally, given the peculiar concepts shown in Fig. 5c we leave the interpretation of this amusing pattern to the reader. We further add that some KBs (such as ConceptNet) contain erroneous, biased and funny facts (Baydin, de Mántaras, and Ontañón 2012) such as these. However, we think that in the context of CC those facts might be fruitful.

Fig. 6 contains three examples of patterns with *ucycles*. The pattern in Fig. 6a represents an entity (*postage stamp*) located somewhere (*post office*) with both the entity and the place having the same activity (*mail letter*) through different possibilities (*usedFor* and *capableOf*) Fig. 6b serves as a good case history of two entities (mathematicians *George Boole* and *Claude Shannon*) having the same *interestfield* but with one influencing the other. Hence, this frame could be classified as the “*source of inspiration*” frame. Lastly, Fig. 6c relates two activities (*dancing* and *playing*) with conditions (*go to party*) as well as outcomes (*having fun*).

Conclusions

We have presented a system designed to discover repeating patterns in large-scale semantic graphs which can be used as frames in computational models of CB. We illustrated how the system mines KBs for interesting patterns using a MOEA and a specialised tool to more efficiently compute the frequency of the patterns involved in the process. We also believe that using frames containing *ucycles* could benefit the blending process with the contribution of closely connected elements with a sense of completeness.

Further Work

We plan to create a repository with the most interesting frames found so far. We could also execute our frame finding system in other KBs such as NELL (Mitchell et al. 2015) to discover even more promising frames. Another idea to be explored is the definition of a large set of useful frames in an ontology. But above all, the KB of frames is expected to be used in a follow-up of our computational model of CB. The impact of frames in the emerging blends will be then better understood, as well as their required characteristics. Those characteristics would then be used to improve our MOEA in the search for suitable frames. We might also improve our *querykb* tool by using a more sophisticated query planning mechanism (for example, one that estimates the sizes of intermediate relations given what is known about the KB). Alternatively, instead of finding the exact number of solutions, we could try to find an approximate count, perhaps by extending the recent algorithm of (Iyer et al. 2018).

Acknowledgements

João Gonçalves is funded by Fundação para a Ciência e Tecnologia (FCT), Portugal, under the PhD grant SFRH/BD/133107/2017.

References

- Baker, C.; Fillmore, C.; and Lowe, J. 1998. The berkeley framenet project. In *Proceedings of the 17th international conference on Computational linguistics-Volume 1*, 86–90. Association for Computational Linguistics.
- Baydin, A. G.; de Mántaras, R. L.; and Ontañón, S. 2012. Automated generation of cross-domain analogies via evolutionary computation. *CoRR*.
- Coulson, S. 2006. Conceptual blending in thought, rhetoric, and ideology. *APPLICATIONS OF COGNITIVE LINGUISTICS* 1:187.
- Cunha, J. M.; Gonçalves, J.; Martins, P.; Machado, P.; and Cardoso, A. 2017. A pig, an angel and a cactus walk into a blender: A descriptive approach to visual blending. In *Proceedings of the Eighth International Conference on Computational Creativity (ICCC 2017)*.
- David, O.; Dodge, E.; Hong, J.; Stickles, E.; and Sweetser, E. 2014. Building the metanet metaphor repository: The natural symbiosis of metaphor analysis and construction grammar. In *The 8th International Conference on Construction Grammar (ICCG 8)*, Osnabrück, Germany.
- Deb, K.; Agrawal, S.; Pratap, A.; and Meyarivan, T. 2000. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii. In *International conference on parallel problem solving from nature*, 849–858. Springer.
- Eppe, M.; Maclean, E.; Confalonieri, R.; Kutz, O.; Schorlemmer, M.; Plaza, E.; and Kühnberger, K.-U. 2018. A computational framework for conceptual blending. *Artificial Intelligence* 256:105–129.
- Fauconnier, G., and Turner, M. 1998. Conceptual integration networks. *Cognitive Science* 22(2):133–187.
- Fauconnier, G., and Turner, M. 2002. *The Way We Think*. New York: Basic Books.
- Fauconnier, G. 1994. *Mental Spaces: Aspects of Meaning Construction in Natural Language*. New York: Cambridge University Press.
- Gangemi, A.; Alam, M.; Asprino, L.; Presutti, V.; and Recupero, D. R. 2016. Framester: a wide coverage linguistic linked data hub. In *European Knowledge Acquisition Workshop*, 239–254. Springer.
- Gonçalves, J.; Martins, P.; and Cardoso, A. 2017. Blend city, blendville. In *Proceedings of the Eighth International Conference on Computational Creativity*.
- Gonçalves, J.; Martins, P.; and Cardoso, A. 2018. Drilling knowledge bases for hidden frames. *Proceedings of TriCoLore 2018*.
- Hadka, D. 2015. Moea framework - a free and open source java framework for multi-objective optimization.
- Iyer, A. P.; Liu, Z.; Jin, X.; Venkataraman, S.; Braverman, V.; and Stoica, I. 2018. {ASAP}: Fast, approximate graph pattern mining at scale. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, 745–761.
- Johnson, C., and Fillmore, C. 2000. The framenet tagset for frame-semantic and syntactic coding of predicate-argument structure. In *1st Meeting of the North American Chapter of the Association for Computational Linguistics*.
- Martins, P.; Pollak, S.; Urbančič, T.; and Cardoso, A. 2016. Optimality principles in computational approaches to conceptual blending: Do we need them (at) all? In *Proceedings of the Seventh International Conference on Computational Creativity (ICCC 2016)*. Paris, France: Sony CSL.
- Martins, P.; Oliveira, H. G.; Gonçalves, J.; Cruz, A.; Cardoso, A.; Znidarsic, M.; Lavrac, N.; Linkola, S.; Toivonen, H.; Hervás, R.; Méndez, G.; and Gervás, P. 2019. Computational creativity infrastructure for online software composition: A conceptual blending use case. *IBM Journal of Research and Development* 63(1):1–17.
- Mitchell, T.; Cohen, W.; Hruschka, E.; Talukdar, P.; Bettebridge, J.; Carlson, A.; Dalvi, B.; Gardner, M.; Kisiel, B.; Krishnamurthy, J.; Lao, N.; Mazaitis, K.; Mohamed, T.; Nakashole, N.; Platanios, E.; Ritter, A.; Samadi, M.; Settles, B.; Wang, R.; Wijaya, D.; Gupta, A.; Chen, X.; Saparov, A.; Greaves, M.; and Welling, J. 2015. Never-ending learning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*.
- Pereira, F. C., and Cardoso, A. 2003. Optimality principles for conceptual blending: A first computational approach. *AISB Journal* 1(4).
- Pereira, F. C. 2005. *Creativity and AI: A Conceptual Blending approach*. Ph.D. Dissertation, University of Coimbra.
- Ruppenhofer, J.; Ellsworth, M.; Petruck, M.; Johnson, C. R.; and Scheffczyk, J. 2016. *FrameNet II: Extended theory and practice*.
- Speer, R., and Havasi, C. 2012. Representing general relational knowledge in conceptnet 5. In *LREC*, 3679–3686.