

# A Large-Scale Computational Model of Conceptual Blending Using Multiple Objective Optimisation

João Gonçalves<sup>\*1</sup>, Pedro Martins<sup>1</sup>, and Amílcar Cardoso<sup>1</sup>

<sup>1</sup>University of Coimbra, CISUC, DEI  
{jcgonc@dei.uc.pt, pjmm@dei.uc.pt, amilcar@dei.uc.pt}

## Abstract

We present a Computational Creativity (CC) system<sup>1</sup> based on Conceptual Blending (CB) theory. To obtain diverse output – in the form of blend spaces – it allows a trade-off in the optimisation process according to multiple criteria. It is handled by a high performance Multi Objective Evolutionary Algorithm (MOEA) supporting a large Knowledge Base (KB), large input spaces, a high number of mappings as well of frames, all in the range of millions of elements in their structures. Some of CB’s Optimality Principles (OPs) were adapted into objectives as well as new ideas we decided to investigate. Initial experiments allows us to conclude that the system exhibits a form of creativity and its output is capable of depict a simple drawing or a short story.

## Introduction

CB theory (Fauconnier and Turner 2002) was proposed as a cognitive theory that explains mechanisms involved in the creation of meaning and insight in the human mind. Those mechanisms are likely to be involved in creation of meaning, argumentation and the communication of thought (Coulson 2006). In the last decade CB has been successfully used in various international computational systems such as (Schorlemmer et al. 2014) and (Martins et al. 2019).

One of the first computational systems to be based on CB was Divago (Pereira 2005) and it is, to the best of our knowledge, the only system to date that thoroughly formalises the OPs and studies their impact on the generated blend spaces.

Divago was followed by a CC system recreated from scratch (Gonçalves, Martins, and Cardoso 2017), with the purpose of going beyond simple toy problems. It achieves that goal by scaling up the quantity of manipulated data: Input Spaces, mappings and frames. In that system the various qualities of the blend space (such as the OPs and other aspects) are combined using predefined weights and evaluated as a single fitness function. However, combining multiple objectives in a single weighted function requires a careful selection of those weights which if incorrectly done can present the user with disparate search spaces (Konak, Coit, and Smith 2006). Additionally, the more objectives there are the more time consuming that selection is.

To solve the above issue our latest computational CB system makes use of Multi Objective Optimisation (MOO). The output of the new system is scattered throughout a Pareto front which contains a diverse set of solutions according to the multiple objectives. Some of these are CB’s Optimality Principles, others are qualities we expect the blend spaces must have. Like its antecessor, the system supports a large KB as the source of Input Spaces (ISs), a high number of mappings and of frames. This allows for a even wider diversity in the system’s output.

Being a Computational Creativity system, we expect our implementation to exhibit some form of **novelty** and **practicality** (Ritchie 2001). The former means that it should create new statements not existing in its KB. The latter that its output should be easy to interpret and be useful, e.g., to serve as a sketch for a drawing or a short story.

The document is organised as follows: after this introductory section we follow with a short description of CB, then with the details of our CB framework and the implemented objectives. It is followed by the section where we justify the usage of a MOO algorithm. Later, we present some of the results, followed by the conclusions of this work and what we expect to do in the future.

## Conceptual Blending Theory

The theory involves interactions between four mental spaces: two *input spaces*, a *generic space* and the *blend space* (Fig. 1), the latter of which contains the product of the CB process. A mental space corresponds to a partial and temporary structure of knowledge built for the purpose of local understanding and action. The ISs supply the content that will be blended. Then, a partial mapping is formed between the ISs. This mapping serves as a sense of similarity (an analogy) between the two spaces and is reflected in the generic space containing the elements shared by the ISs. A selection of this mapping is used to partially project structures from both ISs – including nearby elements – integrating them in an emerging structure called the blend space.

The blend space integrates elements in three sub-tasks: the projection of elements from the ISs into the blend space (**Composition**); using existing knowledge in the form of background *frames* to generate meaningful structures in the blend (**Completion**); executing cognitive work in the blend according to its ongoing emergent structures (**Elaboration**).

<sup>\*</sup>Corresponding author.

<sup>1</sup><https://github.com/jcgonc/blender-mo>




Figure 1: The four mental spaces of CB with two examples of ISs: horse and bird. Best viewed in colour.

CB theory mentions *frames* to guide the blending process towards stable and recognisable wholes. Frames represent situations, stereotypes, interactions or recurring patterns of some sort involving various participants.

The CB process is guided towards highly integrated and coherent blends through eight OPs: **Integration, Topology, Web, Relevance, Unpacking, Intensification of Vital Relations, Maximisation of Vital Relations and Pattern Completion**. The reader is redirected to (Fauconnier and Turner 2002; Geeraerts 2006) to know more about the OPs.

## Implementing Conceptual Blending

To create the two ISs we use a KB of facts given as a set of semantic relations between pairs of concepts. The two ISs are identified using the concept pairs present in the mapping (explained below). The KB and the mental spaces are represented as a semantic graphs of the form  $G(V, E)$  with  $V$  the set of vertices (concepts) and  $E$  the set of ordered (and labelled) pairs of vertices. The  $n$  edges are stored as triples of the form  $(concept_{i,1}, relation_i, concept_{i,2})$ , e.g.  $wing, partOf, bird$ .

Multiple mappings are supported. An individual mapping associates two different subsets of concepts from each IS. A mapping  $M$  is defined as a set of  $n$  ordered concept pairs,  $M = \{p_1 \dots p_n\}, p_i = (concept_{(i,1)}, concept_{(i,2)})$ . Each pair is required to have an order in its components to identify to which IS (1 or 2) does concept component belong to. The first concept always exists in the first IS and the second concept in the other IS. An IS can be identified as the subset of edges of the KB containing either all the first or second concepts of the mapping.

Frames are handled as semantic graphs having variables as its vertices. They are converted to Datalog queries with the purpose of identifying the pattern represented by the frame in the blend space. Each edge of the frame is con-

verted to a Datalog term, the relation to a predicate and each vertice to a *unique* variable. Then all the terms are merged into a conjunctive query. As an example (Fig. 1), the three edges drawn in cyan in the blend space represent a frame converted to the Datalog query:

```
purpose(W,F), ability(H,F), partOf(W,H),
W!=F, W!=H, H!=F.
```

Before querying the Datalog engine the blend space is converted to a set of facts. Then, if the query is satisfied by the facts we conclude that the frame exists in the blend space.

In each epoch the Evolutionary Algorithm (EA) evolves a set of solutions using a mutation operator. Each solution contains in its “chromosome” a blend space and one copy of a mapping chosen randomly from the initially given set of mappings. The mappings are not changed during EA’s execution and each blend space evolves according to the same assigned static mapping.

Solutions are either created when the EA initialises the population or some of the existing solutions are unable to be improved. Then, either the mutation picks up a newly created solution or an existing one from the population. The mutation randomly decides to add an edge to the blend space or to remove an existing edge. The removal is straightforward, giving priority to concepts with a high degree to minimise the fragmentation of the blend space into multiple graph components. The addition follows CB’s “guidelines” by selectively projecting one edge in three possibilities: an edge is simply projected from the KB; an edge is projected from one of the ISs or; an edge is projected from both ISs connecting two concepts of a concept pair (defined in the chromosomes’ mapping). In all three cases the EA chooses the new edge with a high probability of connecting to an existing edge of the blend space. When the projection of the edge uses the mapping one of its concepts is chosen to be: the first concept of the concept pair; the second OR; a blend of both concepts in the form “ $concept_{(i,1)}|concept_{(i,2)}$ ”.

The mutation is followed by the evaluation of the objectives. The solutions are split in  $n$  blocks, given to different executing  $n$  threads and evaluated in parallel. At the same time,  $n$  Datalog engines try to satisfy the frames (converted to queries and previously cached) using the recently generated blend spaces (converted to databases of facts). To evaluate the novelty of the blend space we count the percentage of new edges of the blend not existing in the KB nor in the ISs. An edge is new if it was projected to the blend space using a mapping (i.e., exchanging concepts from a mapping pair or blending them).

## Objectives to Optimise

In this implementation the blend space is not required to satisfy all of CB’s OPs. We implement those we deem appropriate for our intent and create additional objectives to make the blend space have aspects which do not seem to be described in CB theory. The objectives to be optimised are evaluated as real numbers. Depending on the objective, it is either maximised or minimised (the latter is equivalent to the maximisation of its additive inverse). What follows are the implemented objectives and their purpose:




Figure 2: A frame and a blend space integrating the frame.

- (1) **Number of Integrated Frames** - follows CB’s Integration principle. It reflects the number of frames present in the blend space. A frame in the form of a Datalog query is integrated in the blend if the query is satisfied by the blend space after this being converted to a set of statements (see Fig. 2). In our opinion, a blend space with many integrated frames is undesired because it can represent almost any conceptual structure such as a situation, an event, a stereotype, etc. Hence this objective is to be minimised while requiring it to be at least one.
- (2) **Number of Constraints of the Largest Frame** - Fauconnier and Turner’s Pattern Completion principle influences this objective. We interpret a frame as the composition of multiple conditions (semantic relations) such as *X is a Y* and/or *X moves along path Z*. Given a high amount of frames (thousands or more) spanning a diverse number of conditions and a restricted number of semantic relations, a larger frame (according to its number of conditions) will very likely contain the same conditions as a smaller frame (the example two conditions given above are very likely to exist in many frames). By having a higher number of conditions a frame restricts the blend space to a more specific situation or mental structure, something we find more interesting. Therefore from all the frames integrated in the blend, this objective aims to maximise the number of conditions of the largest frame.
- (3) **Mean of Within-Blend Relation Semantic Similarity** tries to bring some variety to our CB’s experiments. We examined a simple application of semantic similarity to the relations present in the blend space by “complementing” the relations of connected pairs of edges. For example, this objective would prioritise a blend space with the relations *partOf* and *ableTo* connected instead of *partOf* linked to another *partOf* (or a *memberOf*). Our reasoning is based on the principle that in this example the former two relations should be more semantically different than the latter two. Given a word embedding containing the relations present in the ISs the algorithm creates a list of semantic cosine similarities (range between -1 and 1) using the two connected relations of all the pairs of connected edges in the blend space; finally it calculates the arithmetic mean of the list. In order to minimise the similarity this objective is also to be minimised.
- (4) **Mean Importance of Vital Relations** - CB mentions certain types of relations - Vital Relations - such as identity, cause-effect, part-whole, space, property, etc. The presence of those relations in the blend is desirable and

they allow for the compression of information present in the ISs. CB describes two OPs related to vital relations - the Maximisation and the Intensification of Vital Relations. Our implementation does not support compression (or intensification) but it does allow for the frequency of some relations to be maximised in the blend space. To achieve this a previously created table mapping an edge label to a weight is used with a higher value meaning the given relation should be more frequent. Then, a weighted sum is calculated using the weight of the relation in the table and its relative frequency in the edges of the blend space. The latter sum is then to be maximised.

- (5) **Input Spaces Balance** prevents the blend space of being the projection of only one (or none) of the ISs. When this objective is maximised the blend contains the same number of concepts projected from the first IS as from the second IS. With  $c_k$  the number of concepts of the  $k = \{1, 2\}$  IS the objective is  $\frac{\min(c_1, c_2)}{\max(c_1, c_2)}$ . This objective has two remarks: when a concept is a blend of two concepts from the first and second ISs, e.g., “surgeon | butcher”, it increases both  $c_1$  and  $c_2$  concept counters; if either  $c_1$  or  $c_2$  are zero the objective is zero to prevent division by zero.
- (6) **Input Spaces Intertwined Mix** - it interconnects concepts of different ISs throughout the emerging blend space, i.e., any concept of one IS should be connected to a concept of the other IS. It is calculated by counting all the edges of the blend space which connect two concepts of different ISs and dividing the count by the number of edges of the blend space. This value is to be maximised.
- (7) **Mean of Words per Concept** - in previous experiments we noticed that the blend space tended to have large concepts such as “*tybee island strand cottage historic district*”, “*south african class exp 2 2-8-0*” or a blend composition of both. Intuitively large concepts are likely to be distinct (the latter example corresponds to a specific railway locomotive). However, larger concepts complicate the interpretation of the blend and therefore an objective was created to help prevent these situations. Hence the addition of an objective to minimise the arithmetic mean of the number of words of each concept in the blend space. The words are counted by splitting the concepts in tokens separated by a white-space character or equivalent (underscore, question mark, etc.).
- (8) **Blended Concepts Ratio** - also noted in previous experiments was that in order to optimise for either the Input Spaces Balance or Input Spaces Intertwined Mix objectives, the blend space tended to be composed of mostly (or only) blended concepts, e.g. “surgeon | butcher” and “scalpel | knife” - an interesting side effect of using an EA. We added a objective to fix this issue by minimising the ratio of blended concepts to the total number of concepts of the blend space. However we agree that it is not a perfect solution as it reduces (and sometimes completely removes) the presence of blended concepts.

## CB as Multi Objective Optimisation

The blend spaces generated by our system must satisfy at least one of objectives. The optimal solutions found by the MOO framework are *Non Dominated* (or *Pareto Optimal*) if none of the objectives can be improved without worsening some of the other. As the MOO framework a MOEA is used having the advantages of relaxing the search for the overall optimum while reducing the time to obtain results we deem good enough. This is of especial importance in our case as we are using a large KB as the source of knowledge. The EA also easily allows the search domain (the blend space) to be a complex structure such as a semantic graph.

## Previous work and Early Experimental Results

The KB supplying the facts was a custom version of ConceptNet V5 (Speer and Havasi 2012) with 1 229 508 concepts, 1 791 604 relations and 39 types of relations <sup>2</sup>.

A set of 1404 mappings was extracted from the KB using (Gonçalves, Martins, and Cardoso 2018) <sup>3</sup>. This mapping framework runs a Genetic Algorithm (GA) to find mappings between the ISs using relation based isomorphisms. On average, each mapping had  $1413 \pm 1188\sigma$  concept pairs.

We used (Gonçalves et al. 2019) <sup>4</sup> to generate the frames. Based on a MOEA, it scholastically generates patterns from existing structures in the KB and checks for the prevalence of those patterns. The idea is that recurrent patterns should represent frames. Resulting patterns are evaluated and selected according to different criteria. Using this framework 76187 frames were extracted from the former KB having on average  $6.3 \pm 1.5\sigma$  edges,  $6.7 \pm 1.4\sigma$  variables and a Relation Semantic Similarity of  $0.34 \pm 0.1\sigma$ .

*epsilon-NSGA-II* was used as the EA. The population size was constant at 2048. Multiple runs were executed where each took the required epochs to reach on average a total of  $12 \pm 6$  hours. The blend spaces were evolved during  $20\,000 \pm 15\,000$  epochs. The solution’s blend spaces were required to have at least 2 integrated frames, an input space balance of at least 10% and the number of concepts between 3 and 10, with this last upper limit set for an easier interpretation of the generated blend space. Otherwise the number of concepts can grow to a large quantity.

A summary of the results is shown in Table 1 and some examples in Fig. 3. In most of the blend spaces relations between typically unrelated concepts can be seen, e.g. *kissing someone is like a mountain*, *an amphibian made of cold blooded-steel* and *a tree made of perforated paper capable of cooling food*, demonstrating objectives 5 and 6 – a intermix of both input spaces, the usage of mappings and the creation of analogies and of simple metaphors. Most blend spaces generated by the system have vital relations (required by the fourth objective): *made of* and *part of* (part-whole), *capable of* (cause-effect) and *isa* (identity) are some of the examples. Objectives 7 and 8 clearly minimised the number

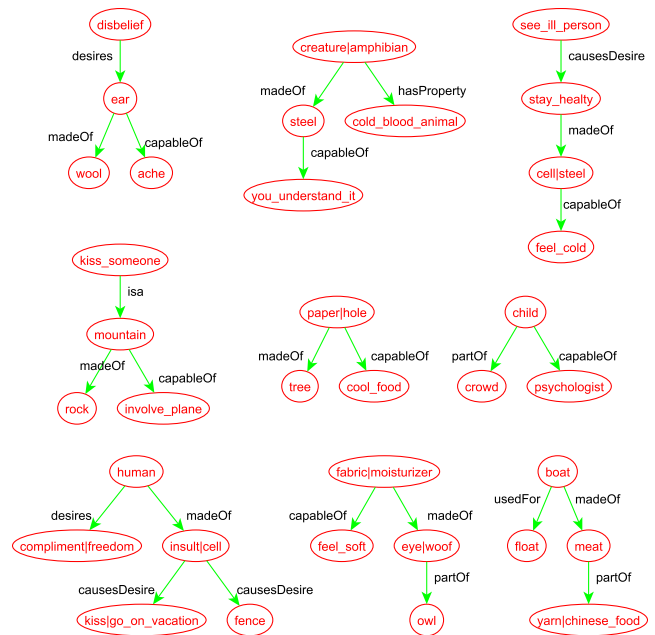


Figure 3: An example of four of the generated blend spaces.

of words in the concepts and the presence of blended concepts in the blend space, respectively. The *Mean of Within-Blend Relation Semantic Similarity* objective brought variety to the presence of somewhat dissimilar connected relations in the blend spaces. It is more likely they contain relations with different labels instead of being dominated by a single relation of the highest importance (because of the fourth objective).

## Further Work

During the implementation and experimentation of this work we noticed several questions. These are the following:

- we think that the way novelty is measured is incomplete as it does not take into account a partial change of the information projected from the ISs. Hence a new way of calculating novelty should be investigated,
- improve the mapping’s quality by incorporating semantic or ontological knowledge between the involved concepts and nearby components. This would allow the system to relate domains of knowledge further disconnected,
- compression of Vital Relations could be implemented through different layers of semantic processing,
- somewhat related to previous question, some form of concept substitution / compression could be applied to concepts related by a relation of semantic similarity or of difference, e.g., replacing a concept by its synonym or antonym. This will likely create further novelty and new relations of interest,
- investigate the EA’s optimisation shortcut addressed by the objective “Blended Concepts Ratio” and a better way of preventing the blend space of being dominated by blended concepts,

<sup>2</sup><https://github.com/jcfgonc/conceptnet-cb>

<sup>3</sup><https://github.com/jcfgonc/ee-mapper>

<sup>4</sup><https://github.com/jcfgonc/pattern-miner>

Objective / Property	Minimum	Mean	Standard Deviation	Maximum
Number of Integrated Frames	0	12.369	15.417	132
Number of Constraints of the Largest Frame	2	4.4356	1.2252	8
Mean of Within-Blend Relation Semantic Similarity	0.03095	0.23884	0.16251	0.89158
Mean Importance of Vital Relations	0.07	0.74376	0.19007	1
Input Spaces Balance	0.14286	0.86858	0.14016	1
Input Spaces Intertwined Mix	0	0.71975	0.22067	1
Mean of Words per Concept	1	1.7680	0.6487	4.875
Blended Concepts Ratio	0	0.28151	0.18553	1
Number of concepts in the blend space	3	5.8738	2.3779	10
Number of edges in the blend space	2	6.5757	4.3176	17
Novelty	0	0.79773	0.20259	1

Table 1: Statistical properties of the optimised objectives, novelty and other blend space’s properties. It contains 7386 solutions.

- instead of requiring strictly equal semantic relations allow the matching of various levels of similar semantic relations in the mapping, frame mining and CB implementation (e.g. matching “partOf” to “component of” or to “belongs to”).

Lastly, a study should be done about how the presence of frames in the blend space influences its quality in some aspect(s) - if positively, negatively or indifferently. Are frames really required or could both the input spaces and mappings contain all knowledge that is required for a good blend?

## Conclusions

We have presented a computational implementation of a MOO system which uses CB theory. The blend’s space criteria in the form of objectives to be optimised were also explained. The system was demonstrated to be capable of handling a large amount of knowledge in the form of input spaces, mappings and frames while at the same time exhibiting creativity in the generated blend spaces.

## Acknowledgements

This work is funded by national funds through the FCT - Foundation for Science and Technology, I.P., within the scope of the project CISUC - UID/CEC/00326/2020 and by European Social Fund, through the Regional Operational Program Centro 2020.

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

## References

- [Coulson 2006] Coulson, S. 2006. Conceptual blending in thought, rhetoric, and ideology. *APPLICATIONS OF COGNITIVE LINGUISTICS* 1:187.
- [Fauconnier and Turner 2002] Fauconnier, G., and Turner, M. 2002. *The Way We Think*. New York: Basic Books.
- [Geeraerts 2006] Geeraerts, D. 2006. *Cognitive linguistics: Basic readings*, volume 34. Walter de Gruyter.
- [Gonçalves et al. 2019] Gonçalves, J.; Bembenek, A.; Martins, P.; and Cardoso, A. 2019. Going into greater depth in the quest for hidden frames. In *Proceedings of the Tenth International Conference on Computational Creativity*.
- [Gonçalves, Martins, and Cardoso 2017] Gonçalves, J.; Martins, P.; and Cardoso, A. 2017. Blend city, blendville. In *Proceedings of the Eighth International Conference on Computational Creativity*.
- [Gonçalves, Martins, and Cardoso 2018] Gonçalves, J.; Martins, P.; and Cardoso, A. 2018. A fast mapper as a foundation for forthcoming conceptual blending experiments. *Special Track Analogy - Proceedings from The Twenty-Sixth International Conference on Case-Based Reasoning*.
- [Konak, Coit, and Smith 2006] Konak, A.; Coit, D. W.; and Smith, A. E. 2006. Multi-objective optimization using genetic algorithms: A tutorial. *Reliability engineering & system safety* 91(9):992–1007.
- [Martins et al. 2019] Martins, P.; Oliveira, H. G.; Gonçalves, J.; Cruz, A.; Cardoso, A.; Žnidaršič, M.; Lavrač, 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).
- [Pereira 2005] Pereira, F. C. 2005. *Creativity and AI: A Conceptual Blending approach*. Ph.D. Dissertation, University of Coimbra.
- [Ritchie 2001] Ritchie, G. 2001. Assessing creativity. In *Proceedings of the AISB’01 Symposium on Artificial Intelligence and Creativity in Arts and Science*, 3–11.
- [Schorlemmer et al. 2014] Schorlemmer, M.; Smail, A.; Kühnberger, K.-U. K.; Kutz, O.; Colton, S.; Cambouropoulos, E.; and Pease, A. 2014. Coinvent: Towards a computational concept invention theory. In *Proceedings of the 5th Int. Conference on Computational Creativity, ICC-14, Ljubljana, Slovenia*.
- [Speer and Havasi 2012] Speer, R., and Havasi, C. 2012. Representing general relational knowledge in conceptnet 5. In *LREC*, 3679–3686.