Optimality Principles for Conceptual Blending: A First Computational Approach

Francisco C. Pereira and Amílcar Cardoso
*Centro de Informática e Sistemas da Universidade de Coimbra (CISUC)
Departamento de Engenharia Informática, Polo II
Pinhal de Marrocos, Coimbra, Portugal
camara@dei.uc.pt, amilcar@dei.uc.pt

Abstract

We propose an implementation of the eight Optimality Principles from the framework of Conceptual Blending, as presented by Fauconnier and Turner (1998). Conceptual Blending explains several cognitive phenomena in the light of the integration of knowledge from different *mental spaces* onto a single *mental space*: the Blend. The Optimality Principles express general pressures that compete in the generation of the Blend.
The work we present now corresponds to the Constraints module of our computational model of Conceptual Blending, also described in other papers.

1 Introduction

One big challenge for Computational Creativity is the generation of new concepts and a very interesting source of inspiration for approaching this issue comes from the framework of Conceptual Blending, from Fauconnier and Turner (1998). The Conceptual Blending (CB) framework establishes cognitive processes and principles that work in the integrations of knowledge needed in reasoning. We are developing a model of computational creativity that takes an approach to CB as a fundamental pilar. This system has been described in earlier papers (Pereira and Cardoso (2001), Pereira and Cardoso (2003)) as well as its first experiments (Pereira and Cardoso (2002)), being evident the need for developing the Optimality Principles of Conceptual Blending as an urgent improvement. These are the general guidelines that drive the process of blending and allow the differentiation between a "good" and a "bad" blend. Here, we propose a formal realization for each of the eight principles.
In the first sections, we give an overview of Conceptual Blending, so as to provide the reader the motivation and background for our approaches to the Optimality Principles. We made some exploratory experiments which are informally reported in the end of the discussion of each principle. We finalize this paper with reflections regarding the creative aspects of this work.

2 Conceptual Blending

Conceptual Blending (CB) was initially proposed by Fauconnier and Turner (1998) as part of a major framework concerning cognition and language and had the role of explaining the integration of knowledge coming from distinct sources onto a single, independent and coherent unit, the Blend. A blend is a concept or web of concepts whose existence and identity, although attached to the pieces of knowledge that participated in its generation (the inputs), conquers gradual independence through time and use.
We find examples of blends in many sorts of situations. A blend can be an effective way to get attention and curiosity towards advertising a product (e.g. Sony's AIBO robot uses all sorts of Sony products behaving as if it were a real human) or spreading a message (e.g. the Marlboro cowboy with impotence problems). People have been making blends with creatures from the times of Greek mythology (e.g. *pegasus*) till today (e.g. *the pokemons*), natural language discourse (e.g. "John digested the book", "Sue sneezed the napkin off the table"), poetry (see Freeman (1999)). It is also claimed to be important in *scientific creativity* (Lakoff and Nunez (2000)) and in the elaboration of *cognitive instruments* (Hutchins (2002)). Many more examples and situations could be listed and studied in detail, demonstrating the ubiquity of CB.
It is noticeable that Conceptual Blending and its research community are growing, and possibly still in its early stages. It is an elegant proposal for a creative process and its relationships with language and cognition, but it carries formal vagueness across its several aspects, making it difficult even to be considered as a theory, in a *Popperian* sense. Indeed, if the big hole between the general model and the specific examples does not show any incoherence, it also does not allow falsifiability, leaving very much undefined the boundary between what is and what is not a blend. These criticisms intend to motivate work that, from our point of view, is fundamental. And support our own motivation for the present project, which is that of contributing with a formal model and implementation

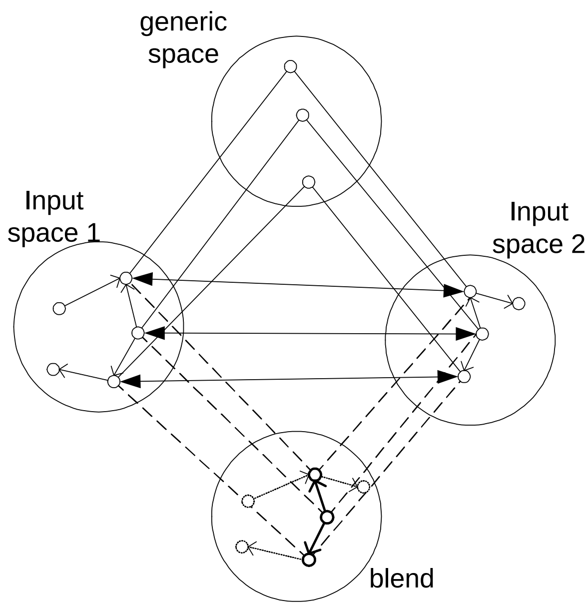

Figure 1: The Conceptual Blending Model

based on Conceptual Blending. We hope it can be useful in shedding some light on the referred issues.

In order to understand the CB framework, we must introduce a fundamental concept: the *mental space*. According to Fauconnier and Turner (1998), mental spaces are partial structures that proliferate when we think and talk, allowing a fine-grained partitioning of our discourse and knowledge structures. For simplifying, let us consider a mental space as a partial selection of knowledge from a domain, a memory of a situation, an imagined scenario or entity, essentially a knowledge structure of inter-related concepts that is explicitly or implicitly necessary for a reasoning. As AI researchers, we see a mental space representable as a *semantic network*, a *frame*, a *case* or any other symbolic knowledge structure that gathers a set of inter-related concepts towards a specific situation.

In its *canonical form*, Conceptual Blending is described as involving two input mental spaces that, recurring to a cross-space mapping between them and a generic space that has general knowledge relevant for both the input domains, will generate a third one, called Blend. This new domain will maintain partial structure from the input domains and have emergent structure of its own. The process of generation of a blend can be summarized according to three general steps (Fauconnier (1997)): Composition, where new relations become available that did not exist in the separate inputs; Completion, when generic knowledge is projected into the blend, to ”complete” the emergent structure; and Elaboration, in which cognitive work is performed in the blend, according to its own emergent logic. The order of these steps may be changed and several iterations of the process may be necessary. There is a set of *governing principles*, the *Optimality Pressures*, that should drive the process of generating a ”good blend”(Fauconnier and Turner (2002)):

- Integration - The blend must constitute a tightly integrated scene that can be manipulated as a unit. More generally, every space in the blend structure should have integration.

- Pattern Completion - Other things being equal, complete elements in the blend by using existing integrated patterns as additional inputs. Other things being equal, use a completing frame that has relations that can be the compressed versions of the important outer-space vital relations between the inputs.

- Topology - For any input space and any element in that space projected into the blend, it is optimal for the relations of the element in the blend to match the relations of its counterpart.

- Maximization of Vital Relations - Other things being equal, maximize the vital relations in the network. In particular, maximize the vital relations in the blended space and reflect them in outer-space vital relations. Turner and Fauconnier identify 15 different vital relations: change, identity, time, space, cause-effect, part-whole, representation, role, analogy, disanalogy, property, similarity, category, intentionality and uniqueness.

- Intensification of Vital Relations - Other things being equal, intensify vital relations.

- Web - Manipulating the blend as a unit must maintain the web of appropriate connections to the input spaces easily and without additional surveillance or computation.

- Unpacking - The blend alone must enable the understander to unpack the blend to reconstruct the inputs, the cross-space mapping, the generic space, and the network of connections between all these spaces

- Relevance - Other things being equal, an element in the blend should have relevance, including relevance for establishing links to other spaces and for running the blend. Conversely, an outer-space relation between the inputs that is important for the purpose of the network should have a corresponding compression in the blend.

As far as we know, there is no work yet towards an objective study of the optimality pressures, measuring examples of blends or specifying these principles in detail. This, we believe, disturbs considerably the appreciation and application of Conceptual Blending in scientific research, making a particular motivation for this work being that of testing and specifying a formal proposal for the optimality pressures.

3 Basic Notions from our model

We present some basic notions that are necessary to understand this document. We use a specific type of mental space, that is static, modeless and non attached to discourse, which is closer to the notion of *domain knowledge* in AI. We call it a *domain*, comprising a theory and a set of instances. In this paper, we consider only the theory, which is represented by a *concept map* that explains the structural and causal organization of the domain. A concept map is a semantic network, with binary directed relations between the concepts[1]. In CB, there is a set of *vital relations* that take a special role in the Blending process. We also consider these relations, but allow the choice of a different set of relations. In principle, these vital relations can be the source for establishing *mappings* between the input spaces, fundamental for the projection operation (two objects mapped to each other can be projected to the same concept, e.g. horse and bird projected to pegasus). Each mapping projection consists on a ternary relation m/3. For example, a mapping algorithm based on "property" could try to map pairs of concepts that make a valid pair of object/property (e.g. "m(property, dark, bird)"), a different one based on analogy could link pairs of analogical counterparts (e.g. "m(analogy, leg, wing)"). As far as our research goes, we are only applying an analogical mapping algorithm that finds a 1-to-1 structure mapping between the concept maps of the two domains. Knowing this is a strong limitation, we hope to address other mapping algorithms in future developments. Currently, for the sake of validation and experimentation of the system, we allow *user-defined* mappings, so as to allow conclusions independent from the mapping choice.

Finally, another important notion is that of *frames*. The reasoning behind a frame lays on the idea that concepts within it should be tightly integrated according to a situation, structure, cause-effect or any other relation that ties a set of concepts onto one, more abstract or broad, composite concept. We envisage different kinds of frames, in terms of level of abstraction. The more specific ones correspond to concepts easily identifiable as familiar objects or situation (which we can see as a kind of *prototypes*). For example, the (very much simplified) frame of "transport means" corresponds to a set of concepts and relations that, when connected together, represent something that has a container and a subpart (e.g. an engine) that serves for locomotion.

$$
\begin{array}{lcl}
\multicolumn{3}{l}{frame(transport\_means(X)):} \\
 & & have(X, container)\wedge \\
carrier(X, people) & \longleftarrow & have(X, Y)\wedge \\
 & & purpose(Y, locomotion)\wedge \\
 & & drive(\_, X)
\end{array}
$$

The more abstract frames can consist on top-level decisions or directives that decide the underlying philosophy of the blend construction. For example, if a blend satisfies the "aframe" frame, it means it maintains the structure of the input domain 1. If it satisfies the "bprojection" frame, then the concepts of input space 2 all get projected unchanged to the blend (e.g. "bird" is projected to "bird", "wing" is projected to "wing"). If a blend satisfies "aframe" and "bprojection" simultaneously, then it should[2] have the concepts from input space 2 organized according to the structure of input space 1. It is more complex to design these more abstract frames, thus we allow the use of programming (in *prolog* language, inside curly brackets) within a frame specification. Below we can see the programming of the "aframe".

$$
\begin{array}{l}
frame(aframe(A)): \\
aframe(A, Blend) \ \longleftarrow \\
\quad \{stats(domain1, A), current\_blend(Bl)\}, \wedge \\
\quad \{findall(R/X/Y, (rel(A, XA, R, YA), pro- \\
\quad jection(Bl, A, XA, X), projection(Bl, A, \\
\quad YA, Y)), L)\} \wedge op(exists(L))
\end{array}
$$

Basically, "aframe" searches for all relations of domain 1 (A) and obtains a list (L) containing their projections to the blend. The $op(exists(L))$ condition is an *operator* that, when, interpreted by the frame processor, expands L into a set of concept map relations (like those in the "transport_means" frame). For "aframe" to be totally satisfied, it is necessary that all relations from input space 1 also exist in the blend.

There are also frames of intermediate level of abstraction, which aren't either as specific as "transport_means" or as abstract as "aframe". For example, "new_creature" is concerned with finding a "creature" (thus having its specific properties - e.g. being a "living being") that didn't exist before in either domains or existed but not as a "creature" (e.g. a "flying snout").

When we say that *a frame f is satisfied in the blend b*, we mean all its premises are true in the domain b. We see frames as *information molds* and building a blend for a given situation should depend much on the choice of these structures, either being structures towards which the blend self-organizes or as pragmatic goals or query specifiers that the blend is expected to accomplish.

[1]First order logic predicates with arity 2, e.g. isa(bird, aves); purpose(wing, fly))

[2]Depending on the mappings, the concepts considered in aframe and bprojection may become separate in the inputs (and so there wouldn't be any systematic relation between concepts from input spaces 1 and 2), yet this would receive little value in the measures presented in this paper.

4 Optimality Principles

Following the F&T notion of Optimality Principles, the pressures that should lead towards stable, integrated new blends, we propose now a set of measures that should reflect as much as possible the rationale behind each principle. In order to give a clearer idea of its individual effect in our blending system, we present a brief report of experiments we made. These experiments consisted in running a parallel search method, a genetic algorithml, to retrieve blends from the search space. The input domains were the domains of *horse* and *bird* (see tables 1 and 2), meaning that the expected results range from the unchanged copy of one (or both) of the concepts to a horse-bird (or bird-horse) which is a combination of selected features from the input domains. The construction of these domains was subject to the following constraints: they should be concept networks in which nodes are concepts and arcs are relations; the concepts should be connected to the ontology in the generic space through an "isa" relation; the relations used should be present in the Generalized Upper Model (GUM) hierarchies (J. Bateman and Fabris (1995)) or be subtypes of them. GUM is a general top-level ontology that has two hierarchies (elements and relations) that comprise abstract relations, properties, spatial relationships, among others. Although allowing a normalization of the concept maps, the constraints in the construction of the domains don't avoid, *per se*, the biasing or ingenious tailoring. For this reason, in this paper and in the exploratory experiments we show, we don't give special attention to a qualitative *reading* of the results or use them to demonstrate its validity, instead we are interested in reporting the effects each measure has on the results and on the search landscape.

The generic domain (in tables 3 and 4) consists on a simple general ontology, a set of frames and integrity constraints. We applied 3 different mappings (figure 2), all generated by the Mapper module. These mappings range from being very small (only four mapping correspondence) to large (21 mapping correspondences), from non-surprising associations (e.g. "animal" and "animal") to nonsense (e.g. snout and lung). For each mapping, we tested the seven optimality pressures. Each of these comprising 30 runs[3].

4.1 Integration

Frames have a natural integration role because they gather knowledge around abstractions, tightening the links between concepts. Assuming the set F of frames that are satisfied in a blend, we define the *frame coverage* of a domain to be the set of relations from its concept map that belong to the set of conditions of one or more frames in F. The larger the frame coverage of the blend, the more it is integrated. Yet, a blend that is covered by many frames should be less integrated than a frame with the same coverage, but with less frames. In other words, if a single frame covers all the relations of a blend, it should be valued with the maximal integration, whereas if it has different frames being satisfied and covering different sets of relations, it should be considered less integrated. The intuition behind this is that the unity around an integrating concept (the frame) reflects the unity of the domain. The Integration measure we propose varies according to this idea. It also takes integrity constraints into account so that, when a frame violates such a constraint, it is subject to penalty.

Definition 4.1 *For a frame f with a set C of conditions $B_i \bigcup \neg B_i$, a blend b, with a concept map CM_b, a blendoid*[4] *CM_{B+}, the concept map of the blendoid and VI, the set of integrity constraints*[5] *that are violated in the frame, the* integration *value, I_f is defined by:*

$$I_f = (\frac{\#C}{\#CM_b} \times (1-\iota)^{\#VI}) \times (1 + \frac{\#CM_b}{\#CM_{B+}})/2$$

being ι an penalty factor between 0 and 1, a value that penalizes a frame for each violation to integrity constraints. An integrity constraint is violated if its premises are true. In the context of the integration measure of frame f above, f violates integrity i if the conditions Ci of i are verified and $Ci \bigcap C \neq \emptyset$. In other words, f needs to violate i in order to be integrated.

We would like to clarify the above formula further more: the first factor represents the ratio of coverage of b w.r.t. f; the second factor means that each integrity constraint violation implies an exponential discount; the third factor serves the purpose of maximizing the size of the blend (if two frames have the same ratio of coverage, the one that contains more relations should have higher integration); the division by 2 aims to normalize the result between 0 and 1.

While the value for a single frame integration is described above, the integration measure of a domain w.r.t. a set of frames is not necessarily straightforward. At a first sight, it is appealing to just sum the values of integration of all frames, or of the union of them. Or even their intersection. But this would lead to wrong results, because a set of frames could not be reduced to a single frame from the point of view of integration. In this measure, we want to stimulate unity, coverage and take into account the strength of each frame individually. In terms of unity, we argue that the set of relations that make the "core" of all the frames that are satisfied, i.e. the intersection of the sets C of conditions of all frames, should be highly valued. On the other side, the coverage of this "core" will be smaller than the overall coverage (or equal, if

[3]A run is an entire evolutive cycle, from the initial population to the population in which the algorithm stopped

[4]The concept map that contains all the possible relations the blend may have

[5]Rules with *false* conclusion as in table 3

the frames have equivalent C sets), which leads us to take into account the disjoint sets of relations of the frames. Finally, the integration of each individual frame (as defined above) should also be present in the overall measure. These last two issues (the overall coverage and the integration of individual frames) are subject to a disintegration factor because they reflect the existent of different, not totally intersected, frames. We propose this factor, α, to be a configurable value from the interval $[1, 0]$. It is now time to present our proposal for the $IntegrationMeasure$ of a blend:

Definition 4.2 *For a set of frames $f_i \in F_b$, with F_b being the set of the frames that have their conditions (C_i) satisfied in the blend b*

$$Integration = I_{\bigcap_0^i C_i} + \alpha \times Uncoverage \times \sum_0^i I_{f_i}$$

The $Uncoverage$ value consists on the ratio of relations that do not belong to the intersection of all frames w.r.t. the total number of relations considered in the frames:

$$Uncoverage = \frac{\# \bigcup_0^i C_i - \# \bigcap_0^i C_i}{\# \bigcup_0^i C_i}$$

We think the integration measure is a fundamental brick of the blending process. It leads the choice of the blend to something *recognizable* as a whole, fitting patterns that help to determine and understand what a *new* concept is.

Experiments: An immediate conclusion about the effect of Integration is that frames behave as *attractor* points in the search space. Moreover, the frames with a larger coverage tend to be preferred, although when too large (like *aprojection* or *aframe*) they are dropped away. The evolution is directed to a compromise of coverage and satisfiability. More specifically, when it satisfies a frame like *pw_based_explanation*, the resulting Integration value is a local maximum or a point in its neighborhood (because sometimes other, related, frames were also found) and ”jumping” to another area of the search space becomes difficult.

Another conclusion to take concerns to the observation that the complexity of the search space landscape grows with mapping size. In fact, when we have a mapping of size 2, the algorithm only finds two different solutions and the better rated (possibly a global maximum) is achieved in 77% of the runs, but with a mapping of size 5, it returns six different blends, being the best choice retrieved only 43% of the times. To confirm this conclusion, the mapping of size 21 lead the algorithm to 16 different maxima, being the best one found only 7% of the times. A good compensation for this apparent loss of control is that the returned values are clearly higher (0.68, for the best) than in the small mappings (0.22), meaning that with big mappings there are many more possibilities to find integrated blends.

The resulting concept maps consist on exactly the relations that are covered by the satisfied frame or combination of them, more specifically there were two frames that were very persistent: *pw_based_explanation* and *purposeful_subpart*. In some blends, the former was present multiple times (e.g. a part-whole explanation of a horse, with part-whole explanations of its subparts) and both were plenty of times combined.

4.2 Topology

The Topology optimality pressure brings *inertia* to the blending process. It is the constraint that drives against change in the concepts because, in order to maintain the same topological configuration as in the inputs, the blend should maintain exactly the same neighborhood relationships between every concept, ending up being a projected copy of the inputs. In *real blends*, this pressure is normally one that is disrespected without big loss in the value of the blend. This is due to the *imagination* context that normally involves blends, i.e., novel associations are more tolerable.

In our Topology measure, we follow the principle that, if a pair of concepts, x and y, is associated in the blend by a relation r, then the same relation must exist in the inputs between the elements from which x and y were projected. We say that $r(x, y)$ is *topologically correct*. Thus, the value of Topology corresponds to the ratio of topologically correct relations in the concept map of the blend.

Definition 4.3 *For a set $TC \subseteq CM_b$ of* topologically correct *relations, defined as*

$$TC = \{r(x, y) : r(x, y) \in CM_1 \cup CM_2)\}$$

where CM_1 and CM_2 correspond to the concept maps of inputs 1 and 2, respectively[6]. *The topology measure is calculated by the ratio:*

$$Topology = \frac{\#TC}{\#CM_b}$$

Intuitively, this measure represents the amount of relations from the inputs that got projected onto the topologically equivalent position in the blend. At the moment, the only way to violate topology is by having a pair of concepts projected to the same one (e.g. ”horse” and ”bird” projected to ”horse”), bringing a new relation exclusive to one the domains (e.g. ability(bird, fly) projects to ability(horse, fly)). Topology thus decreases as fusion or transfer projections are made.

Experiments: In all the experiments with Topology, the final results were valued 100%, meaning that this constraint is easily fully accomplished, independently of the

[6] In other words, TC is the intersection of the concepts maps of the blend and the input spaces

mapping. An interesting fact is that there is a multitude of solutions in the *search landscape* of Topology, showed by the amount of different final results in each mapping. Intuitively, and observing the short duration of each run, this means that, wherever the search starts, there is always a Topology optimal point in the neighborhood.

Topology is more an inertia than a transformation force because it values knowledge that remains the same. In our horse-birds, this pressure projects wings, beaks and claws (i.e. concepts from the bird domain) to the blend but isolates them unless there is strong evidence to connect to horses, legs and snouts (i.e. concepts from the horse domain).

4.3 Pattern Completion

The Pattern Completion pressure brings the influence of patterns being them present in the *inputs* or come from the *generic* space. Sometimes a concept (or a set of concepts) may seem incomplete but making sense when ”matched against” a pattern.

At present, in the context of this work, a pattern is described by a frame, i.e. we don't distinguish these two concepts, and therefore pattern completion is basically frame completion. Here, as in the definition of this principle, the completing knowledge becomes available from ”outside”, not as a result of projection. This means the act of completing a frame consists on asserting the truth of the ungrounded premises, a process that happens only after a sufficient number of premises is true. We call this the *evidence threshold*. An interesting approach in Pattern Completion we would like to consider is that of *abduction*. I.e., if the conclusions are satisfied, why not derive the premises?

The evidence threshold e of a frame f_i with regard to a blend b is calculated according to the following.

$$e(f_i, b) = \frac{\#Sat_i}{\#C_i} \times (1 - \iota)^{\#VI}$$

where Sat_i contains the conditions of each f_i that are satisfied in b, ι is the integrity constraint violation factor and VI the set of violated integrity constraints.

As in the integration pressure, we have the problem of taking into account multiple frames. This time, given that we are evaluating possible completion of subsets of relations, instead of sets of relations that are actually verified in the domain, it is difficult to find such a linear rationale (e.g. would two patterns each with individual completion x value higher than three having each slightly less than x?). As a result, we propose to find the union of the patterns and then estimate its own evidence threshold:

Definition 4.4 *The Pattern Completion measure of a blend b with regard to a set F with n frames is calculated by*

$$PatternCompletion = e(\cup_0^n f_i, b)$$

This measure has a very important role in increasing the potential of the blend, for it brings the ”seeds” that may be used in the Completion and Elaboration phases.

Experiments: The first conclusion to take from the experiments with Pattern Completion is that the size of the resulting concept maps tend to grow as the evolution progresses, although there is no linear correlation found. This has a simple interpretation, given that this measure stimulates the appearance of patterns (frames) that are only partially completed. In doing so, it drives the blend to partially complete (i.e., instantiate partially its conditions) the highest possible number of frames, leading, in each case, to several sets of relations that fit into those frames without satisfying them (e.g. wings are projected, they serve to fly, but they are not attached to anything).

In which respects to the *search landscape*, it seems to be very rich in local maxima. This fact is not unexpected, considering the discussion of the previous paragraphs, the number of different frames available (9) and all their different combinations. The most constant results came from mapping 2, with 13% of the best result obtained and 20% of the second best. An interesting remark is that the resulting local maxima always fall within a very strict range of values (of maximum amplitude 0.11).

4.4 Maximization of Vital Relations

For the maximization of vital relations, we estimate the impact of the vital relations to the blend calculated by the ratio of vital relations w.r.t. the whole set of possible vital relations the blendoid. The blendoid is the largest possible blend that can be obtained from a given mapping and is calculated by projecting to it every concept (i.e., there is no selective projection) regardless of integrity constraint violation or any other constraint. Since it has the largest set of potential relations, it also has the maximum possible of vital relations.

Definition 4.5 *Let Υ be a set of* vital relations*. From the concept map of the blend b, we may obtain the set of vital relations in b, BVR:*

$$BVR = \{r(x, y) : r(x, y) \in CM_b \wedge r \in \Upsilon\}$$

From the blendoid (the largest possible blend), B^+, we have BVR^+:

$$BVR^+ = \{r(x, y) : r(x, y) \in CM_B^+ \wedge r \in \Upsilon\}$$

Finally, the Maximization of Vital Relations measure is calculated by the ratio

$$Maximization\_VR = \frac{\#BVR}{\#BVR^+}$$

Experiments: The influence of Maximization of Vital Relations in the results is straightforward, given that its highest value (1) reflects the presence, in the blend, of all the vital relations that exist in the inputs, independently of the projections of the concepts or non-vital relations (which become *noise* in the sense that these appear randomly and making no difference to the value of the measure, yet confusing the ”reading” of the concept map). As the evolution goes on in each run, the value grows until reaching the maximum reasonably early. For each set of 30 runs, it reached the value 1 a minimum of 93% of the times, and the remaining achieved at least a value of 0.95. As in Topology, the search space of Maximization of Vital Relations is very *simple* since there is a global maximum in the neighborhood of (almost) every point. On the other hand, since there is control on *noise*, the resulting concept maps show unity only in the subset of vital relations.

4.5 Intensification of Vital Relations

Intensification of Vital Relations is the principle that maximizes the concentration around a specific vital relation. I.e., while the Maximization of Vital Relations favors the creation in the blend of vital relations in general as opposed to ”regular” relations, Intensification is based on focussing a specific vital relation. The former relates ”new” vital relations with ”new” relations in the blend; the latter relates vital relations with themselves. Thus, we need a notion of ”intensity” of a vital relation. For such, we argue that a vital relation is considered more ”intense” when there is more evidence of its strength. This evidence should be dependent on the kind of vital relation we are dealing with. For example, an ”analogy” vital relation between two concepts is stronger when there is also a systematical association between the neighborhood concepts (the systematicity principle). In fact, systematicity is the only ”intensity” heuristic we have now and its calculation ($Int_{analogy}$) is straightforward. For a mapping (of size n)

$$Int_{analogy} = \frac{\#analogical\_transfers}{n}$$

where an *analogical transfer* consists on a projection of a concept x from input space 1 to y in the blend (where y is the analogical counterpart of x in the input space 2).
Considering several different ”intensity” heuristics, the evaluation of this pressure takes the point of view that a blend that applies only one vital relation, with intensity x, should have higher measure than a blend with n vital relations, each with intensity x/n (the sum would thus ne x). So we want to favor ”concentration” of vital relations.

Definition 4.6 *Let $\upsilon_i \in \Upsilon$ be a vital relation and the set VRB_{υ_i}, of the instances of vital relation υ_i in the blend, defined by $VRB_{\upsilon_i} = \{\upsilon_i(x,y) : \upsilon_i(x,y) \in CM_b\}$*
Assuming a value Int_{υ_i} of intensity of the vital relation υ_i, the measure of Intensification of Vital Relations is calculated by:

$$Intensification\_VR = \frac{\sum_0^n Int_{\upsilon_i}^2}{(\sum_0^n Int_{\upsilon_i})^2},\ n = \#\Upsilon$$

Intensification is thus higher when there is more ”concentration” (e.g. $Int_{V_1} = 2$, $Int_{V_2} = 2 \Rightarrow Intensification = 8/16$; $Int_{V_1} = 4 \Rightarrow Intensification = 1$).
Experiments: The behavior of Intensification of Vital Relations is similar to that of Maximization in the sense that the search landscape is not very complex. From wherever the search starts, there is a high probability (of at least 70%, in the worst case) of finding the maximum value in the neighborhood. It is important to point, though, that we are only applying one heuristic (analogical transfer), and so the results couldn't be different. Given this fact, we must say that we cannot claim or discuss much about this Intensification of V.R. proposal unless we find other heuristics and mapping procedures other than analogy.

4.6 Unpacking

Unpacking is the ability to reconstruct the whole process starting from the blend. In our view, such achievement underlies the ability to reconstruct the input spaces, specifically. I.e., the reconstruction of the input spaces from the blend demands the assessment of the cross-space mappings, the generic space and other connections. Thus, what we are proposing is that Unpacking can be reduced to the ability to reconstruct the inputs. This is so because there is no way to properly reconstruct the inputs without a reconstruction of the cross-space mappings, generic space and the connections between spaces.
Unpacking should take the point of view of the ”blend reader”, i.e., someone or something that is not aware of the process of generation, thus not having access to the actual projections. Being such, this ”reader” will look for patterns that point to the ”original” concepts. Once again we use the idea of *frames*, more specifically the *defining frame* of a concept, which comprises its immediately surrounding concepts and relations. For example, if the concept ”wing” were projected onto x in blend, the defining frame with regard to the ”bird” domain would consist on *purpose(x, fly), conditional(x, fly), quantity(x, 2)* and *pw(x, bird)*. The more of these relations are found in the blend, the more likely it is that the ”reader” will find easy to understand the relationship between x and ”wing”.

Definition 4.7 *Given a blend b and an input space d, the concept x (which is the projection of the element x^d of input space d to the b), has a defining frame $f_{x,d}$ in d consisting on*

$$f_{x,d} = C_0, C_1 \dots C_n \longrightarrow true$$

where $C_i \in \{r(x,y) : r(x^d,y) \in CM_d\}$.

Assuming that k is the number of conditions (C_i) of $f_{x,d}$ that are satisfied in the blend, the unpacking value of x with regard to d (represented as $\xi(x,d)$) is

$$\xi(x,d) = \frac{k}{n}$$

We calculate the *total estimated unpacking value* of x as being the average of the unpacking values with regard to the input spaces. Thus, having input spaces 1 and 2, we have

$$\xi(x) = \frac{\xi(x,1) + \xi(x,2)}{2}$$

Definition 4.8 *Let $\mathcal{X}$ be the set of n concepts of the blend b, generated from input spaces 1 and 2. The Unpacking value of b is calculated by*

$$Unpacking = \frac{\sum_{i=0}^{n} \xi(x_i)}{n}, x_i \in \mathcal{X}$$

Experiments: The results of the Unpacking measure show that it has a notorious side effect on the size of the blend, it drives it to very small sets (between 0 and 5) of relations. The interpretation here is straightforward: the ratio of *unpackable* concepts is highly penalized in bigger sets because of the projected relations that come as side effect of the projection of (*unpackable* or not) concepts. These relations *confuse* the unpacking algorithm so that it leads the evolution to gradually select the smaller results.
The maxima points also correspond to the value 1, but it seems, from the experiments, that there is a very limited set of such individuals, achieved in the majority (at least 93% for each mapping) of the experiments.

4.7 Web

The Web principle concerns to being able to ”run” the blend without cutting the connections to the inputs. In our opinion, this is not an independent principle, being co-related to those of Topology and Unpacking because the former brings a straightforward way to ”maintain the web of appropriate connections to the input spaces easily and without additional surveillance or computation” and the latter measures exactly the work needed to reconstruct the inputs from the blend. It is not to say that Web is the same as Topology or Unpacking, what we are arguing is that, on one side, Topology provides a pressure to maintain the most fundamental connection to the input: the same structure; on the other side, Unpacking evaluates the easiness of reestablishing the links to the inputs. These two values combined in a weighted sum yield, we propose, an estimation of the strength of the web of connections to the inputs:

$$Web = \alpha \times Topology + \beta \times Unpacking$$

with $\alpha + \beta = 1$. Since this is not an independent variable, making independent experiments with the Web measure would not add any valuable conclusion. In a subsequent publication, we plan to focus on correlation of measures, where we may explore the behavior of this measure.

4.8 Relevance

The notion of ”relevance” or ”good reason” for a blend is tied to the context and goal of the blending generation. A blend, or a part of it, may be more or less relevant dependent of what it is for. Once again, frames take a fundamental role as being ”context” specifiers, (i.e., the set of constraints within a frame describe the context upon which the frame is fulfilled). Therefore, having a set of goal frames, which could be selected from any of the existent domains or specified externally, a blend gets the maximum Relevance value if it is able to satisfy all of them.
An aspect of the goal frames is that they allow the application of queries. For example, if we want to find a concept that ”flies”, we could build a goal frame with the relation $ability(x, fly)$. The blends that satisfy this frame would have high relevance.

Definition 4.9 *Assuming a set of goal frames, F_g, the set F_b of the satisfied frames of blend b and the value PCN_F for the pattern completion of a set of frames F in blend b, as described in section 4.3, we have*

$$Relevance = \frac{\#(F_g \cap F_b) + \#F_u \times PCN_{F_u}}{\#F_g}$$

where F_u, the set of unsatisfied goal frames, consists on $F_u = F_g - F_b$.
Experiments: The first part of the test on Relevance focussed on making a single relation query. In this case, we asked for ”something that flies” (ability(_, fly)). The results were straightforward in any mapping, accomplishing the maximum value (1) in 100% of the runs, although the resulting concept maps did not reveal necessarily any overall *quality* or unity. In other words, the evolution took only two steps: when no individual has a relation ”ability(_, fly)”, therefore with value 0; when a relation ”ability(_,fly)” is found, yielding a value 1, independently of the rest of the concept map.
The second part of the test on Relevance, by adding a frame (*ability_explanation*) to the query, revealed similar conclusions. There was no sufficient knowledge in any of the input domains to satisfy this new frame completely, so the algorithm searched for the maximum satisfaction and reached it 100% of times in every mapping. So the *landscape* seems to have one single global and no local maxima, reflecting the integration of the two parts of the query. If there were separate frames, it is expectable the existence of local maxima. Intuitively, the *search landscapes* of Integration and Relevance seem to be similar.

5 Discussion and further work

A first conclusion we took from this work says that the eight optimality principles can be reduced to seven (since Web is not independent). Even more, given the power of language we use in frames, some of the principles can be *coded* within a frame, namely Topology and Unpacking, and accomplished via a query measured by

Relevance. This reduces our number to five. Yet, we don't know whether this reduction reflects a lack in the CB framework or in our interpretation of it.

We have already stated that the main motivation of our system is to generate new concepts out of previous knowledge. An interesting study to follow, and which has already been made (Pereira and Cardoso (2003)), concerns the necessary parameters that lead to specific concepts (e.g. a pegasus) and how do the changing on these bring different concepts (other creatures). Such experiments provide a bigger insight on the creative potential of our system. Perhaps a more important study, in the sense that it brings a validation to our model, is the application of literature examples from CB (Coulson (2000)) and Conceptual Combination (Costello and Keane (2000)), observing the effects of the application of optimality constraints presented in this paper as well as frames.

6 Acknowledgements

We would like to thank the reviewers for their constructive comments and to Pablo Gervás and Margaret Freeman for participating in discussions that gave rise to some ideas in this paper.

References

Fintan J. Costello and Mark T. Keane. Efficient creativity: Constraint-guided conceptual combination. *Cognitive Science*, 2000.

Seana Coulson. *Semantic Leaps: Frame-shifting and Conceptual Blending in Meaning Construction*. New York and Cambridge: Cambridge University Press, 2000.

Gilles Fauconnier. *Mappings in Thought and Language*. Cambridge University Press, 1997.

Gilles Fauconnier and Mark Turner. Conceptual integration networks. *Cognitive Science*, 22(2):133–187, 1998.

Gilles Fauconnier and Mark Turner. *The Way We Think*. Basic Books, 2002.

Margaret Freeman. The role of blending in an empirical study of literary analysis. In *Proceedings of the 6th International Cognitive Linguistics Conference*, 1999.

Edwin Hutchins. Material anchors for conceptual blends. *forthcoming*, 2002.

B. Magnini J. Bateman and G. Fabris. The generalized upper model knowledge base: Organization and use. *Towards very large knowledge bases: knowledge building and knowledge sharing*, 1995.

George Lakoff and Rafael Nunez. *Where Mathematics Comes From: How the Embodied Mind brings Mathematics into Being*. New York:Basic Books, 2000.

Francisco C. Pereira and Amilcar Cardoso. Knowledge integration with conceptual blending. In *Proceedings of AICS 2001*, 2001.

Francisco C. Pereira and Amilcar Cardoso. The boat-house visual blending experience. In *Proceedings of the Symposium for Creativity in Arts and Science of AISB 2002*. AISB, 2002.

Francisco C. Pereira and Amílcar Cardoso. The horse-bird creature generation experiment. *submitted*, 2003.

ear	$\leftrightarrow$	wing
snout	$\leftrightarrow$	bird
eye	$\leftrightarrow$	lung
mouth	$\leftrightarrow$	feathers
2	$\leftrightarrow$	2
hear	$\leftrightarrow$	fly

1

mouth	$\leftrightarrow$	beak
snout	$\leftrightarrow$	bird
eye	$\leftrightarrow$	lung
ear	$\leftrightarrow$	feathers
eat	$\leftrightarrow$	eat

2

vegetable_food	$\leftrightarrow$	vegetable
food	$\leftrightarrow$	food
horse	$\leftrightarrow$	bird
equidean	$\leftrightarrow$	aves
animal	$\leftrightarrow$	animal
human_setting	$\leftrightarrow$	house
wilderness	$\leftrightarrow$	wilderness
ruminant	$\leftrightarrow$	oviparous
run	$\leftrightarrow$	fly
cargo	$\leftrightarrow$	pet
neigh	$\leftrightarrow$	chirp
snout	$\leftrightarrow$	lung
mane	$\leftrightarrow$	feathers
tail	$\leftrightarrow$	beak
leg	$\leftrightarrow$	eye
hoof	$\leftrightarrow$	wing
4	$\leftrightarrow$	2
eye	$\leftrightarrow$	leg
ear	$\leftrightarrow$	claw
hear	$\leftrightarrow$	catch
grass	$\leftrightarrow$	grass

3

Figure 2: The three mappings

isa(horse,equinae)	pw(leg, horse)	purpose(horse, food)
isa(equinae,mammal)	purpose(leg, stand)	sound(horse, neigh)
existence(horse, farm)	pw(hoof, leg)	purpose(mouth, eat)
existence(horse, wilderness)	purpose(horse, traction)	purpose(ear, hear)
pw(snout, horse)	eat(horse, grass)	color(mane, dark)
pw(mane, horse)	ability(horse, run)	size(mane, long)
pw(tail, horse)	carrier(horse, human)	material(mane, hair)
quantity(hoof, 4)	quantity(leg, 4)	
pw(eye, snout)	quantity(eye, 2)	
pw(ear, snout)	quantity(ear, 2)	
pw(mouth,snout)	purpose(eye, see)	
motion_process(horse,walk)	ride(human, horse)	
taxonomicq(horse, ruminant)	purpose(horse, cargo)	

Table 1: The domain theory of *horse*

isa(bird,aves)	existence(bird, house)	isa(parrot, bird)
isa(aves,oviparous)	purpose(bird, pet)	isa(nest, container)
lay(oviparous, egg)	existence(bird, wilderness)	role_playing(bird, freedom)
purpose(bird, food)	purpose(eye, see)	ability(parrot, speak)
smaller_than(bird, human)	purpose(beak, chirp)	purpose(claw, catch)
pw(lung, bird)	motion_process(bird, fly)	purpose(wing, fly)
purpose(lung, breathe)	quantity(eye, 2)	pw(claw, leg)
isa(paradise_bird, bird)	quantity(wing, 2)	pw(beak, bird)
isa(owl, bird)	quantity(claw, 2)	pw(eye, bird)
ability(bird, fly)	pw(wing, bird)	quantity(leg, 2)
pw(feathers, bird)	conditional(wing, fly)	pw(leg, bird)
purpose(beak, eat)	sound(bird, chirp)	purpose(leg, stand)
pw(straw, nest)		

Table 2: The domain theory of *bird*

isa(entity,something)	isa(mammal, animal)	isa(house, human_setting)
isa(situation,something)	isa(animal, living_entity)	isa(setting,space_location)
isa(state,situation)	isa(oviparous, animal)	isa(space_location,spatial_entity)
isa(process,situation)	isa(behavior, property)	isa(physical_entity,spatial_entity)
isa(temporal_entity,entity)	isa(human_setting, setting)	isa(physical_object,physical_entity)
isa(spatial_entity,entity)	isa(bird, existence, wilderness)	isa(property_measure,information_entity)
isa(information_entity,entity)	isa(wilderness, setting)	isa(imaginary_spatial_entity,spatial_entity)
isa(human, primate)	isa(farm, human_setting)	isa(property,information_entity)
isa(primate, mammal)	isa(equinae, mammal)	isa(living_entity,physical_entity)

⋮

shape(X, Y),shape(X,Z), Y$\neq$ Z $\rightarrow$ false
quantity(X, Y), quantity(X,Z), Y$\neq$ Z $\rightarrow$ false
behavior(X, friendly), behavior(X, dangerous)$\rightarrow$ false
actor(X,_), not isaN(X, action)$\rightarrow$ false
pw(A,A)$\rightarrow$ false

Table 3: The generic space concept map and integrity constraints

Frame name	**Conditions**
aframe	The blend contains the same relational structure of input 1
aprojection	The blend contains the same concepts of input 1
bframe	The blend contains the same relational structure of input 2
bprojection	The blend contains the same concepts of input 2
pw_based_explanation	Set of part-whole relations associated to a concept
transport_means	Features expected in a generic transport means
purposeful_subpart	Set of relations that justify the existence of a subpart of a concept
new_ability	A concept has an ability relation not existent in any of the inputs
new_creature	A concept is a living thing that did not exist (or wasn't such) in any of the inputs
new_feature	A concept has a feature relation not existent in any of the inputs
ability_explanation	A concept has an ability and it is explained by several features (*what does the ability*, *what is it for*, *any conditions necessary*)

Table 4: Frames of the generic space