The Good, the Bad, and the AHA! Blends

P. Martins[1], T. Urbančič[2,3], S. Pollak[3], N. Lavrač[3,2], A. Cardoso[1]
[1] CISUC, Department of Informatics Engineering, University of Coimbra, Coimbra, Portugal
[2] University of Nova Gorica, Nova Gorica, Slovenia
[3] Jožef Stefan Institute, Ljubljana, Slovenia

Abstract

We present and discuss quality assessment of visual blends based on how humans perceive them. This work represents part of a wider study aimed at determining the fundamental characteristics of a good blend. Based on the obtained insights, we hope to make a more comprehensible explanation of some less clear and not fully described aspects of the conceptual blending mechanism that play a fundamental role in creative thinking. Additionally, we intend to bring these insights into the design of artificial creative systems.

Introduction

Conceptual blending (CB) is a vital cognitive mechanism by which two or more mental spaces are integrated to produce new concepts (Fauconnier and Turner 2002). Blending is at the heart of the origin of ideas; a new idea or thought can be seen as an insight gathered from a *blend*, i.e., the result from integrating different mental spaces. Not unexpectedly, the complexity and the quality of blends can be quite heterogeneous. The human brain continuously attempts to blend different concepts either by using a quite uncomplicated web of mental spaces or a more refined and complex network. The majority of these attempts fail in producing good blends, especially because the blend neither has sufficient novelty nor it has immediate purpose (Turner 2014).

What makes us prefer one blend to another? Is it sufficient to require novelty and value under a given context? Can quality simply depend on the coherence of the blend and on the easiness to interpret it? As CB theory inevitably links with the phenomenon of creativity, that is, the ability of producing new, surprising, and valuable ideas or artifacts (Boden 1991), it is expectable to regard as good blends the ones which imply a more creative thinking. However, the "intuitive" nature of creativity hinders the construction of a system of strict and immediate rules to explain which mental spaces should be selected and how they should be integrated in order to achieve creative thinking. Therefore, giving a more elaborate answer to the question "What makes a good blend?" is challenging.

In this paper, we present and discuss quality assessment of visual blends based on how humans perceive them – as *good*, as *bad*, or as *surprising* and *thought-provoking* (*AHA!*) blends. This is part of a wider study whose fundamental goal is to understand what are the key characteristics of a good blend. By finding them, we hope to make a more understandable explanation of some less clear and less described aspects of the blending mechanism that play a fundamental role in creative thinking. Another goal of our work is to bring these insights into the design of artificial creative systems to improve creation and curation processes, especially when they rely on computational models of conceptual blending. We are particularly interested in contributing to the design of frameworks for creativity assessment. While there are already noteworthy works in the field (Ritchie 2001; Colton, Pease, and Ritchie 2001; Wiggins 2001; Colton 2008; Jordanous 2012), there is still room for improvement.

To the best of our knowledge, an analysis of blends based on human perception was only followed by Joy et al. when analyzing conceptual blending in advertising (Joy, F. Sherry Jr., and Deschenes 2009). The authors conducted interviews with 28 volunteers who had to interpret several advertisements by describing what they thought was their main messages and how they arrived at such interpretation. The advertisements used in the experiment provided clear examples of conceptual blending.

We make use of an online-survey questionnaire in which participants are asked to evaluate criteria that we assume to be related to the quality of blends. In the likeness of the aforementioned work, the examples given to the participants can be easily perceived as instances of conceptual blending.

The remainder of this paper is organized as follows. In the upcoming section, we overview the conceptual blending framework and discuss its relation with creativity. Then, we present the content of the survey and discuss its results. Finally, we present concluding remarks and discuss future research.

Conceptual Blending and Creativity

Fauconnier and Turner originally proposed conceptual blending theory as an attempt to explain cognitive and linguistic phenomena such as metaphor, metonymy, and counterfactual reasoning (Fauconnier and Turner 1998), but later it was extended to describe and explain different cognitive phenomena related to the creation of ideas and meanings (Fauconnier and Turner 2002; Turner 2014).

Mental spaces network

A key element in conceptual blending is the *mental space*, which corresponds to a partial and temporary knowledge structure created for the purpose of local understanding (Fauconnier 1994). Mental spaces differ from *frames*, which are more stable knowledge structures. In the CB framework, there is a network comprising at least four connected mental spaces, as depicted in Figure 1. Two or more of them correspond to the *input spaces*, which are the initial domains. A partial matching between the input spaces is constructed. This association is reflected in another mental space, the *generic space*, which contains elements common to the different input spaces. The latter space captures the conceptual structure that is shared by the input spaces. The outcome of the blending process is the *blend*, a mental space that maintains partial structures from the input spaces combined with an emergent structure of its own.

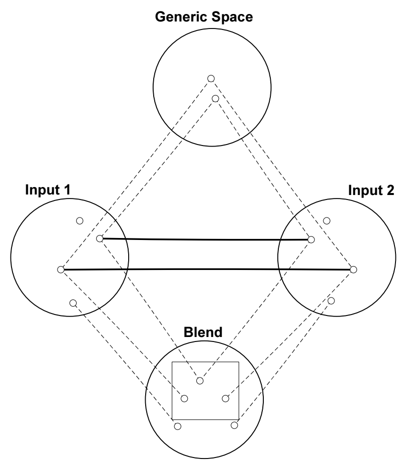

Figure 1: The original four-space conceptual blending network (Fauconnier and Turner 2002).

Integration

Integration of input elements in the blend space results from three operations: *composition*, *completion*, and *elaboration* (Fauconnier and Turner 2002). Composition occurs when the elements from the input spaces are projected into the blend and new relations become available in the blended space. This implies projecting into the blend not only the matched elements but also other surrounding elements. Completion occurs when existing knowledge in long-term memory, i.e., knowledge from background frames, is used to generate meaningful structures in the blend. Elaboration is an operation closely related to completion; it involves cognitive work to perform a simulation of the blended space. Elaboration is also known as "running the blend". There is not a pre-established order for these operations and several iterations may occur.

Optimality principles

Integration is guided by *optimality principles*, which are responsible for generating consistent blends which in turn are more easily interpreted. Fauconnier and Turner (1998) provided a list of these principles:

OP1 *Integration*: the blend must constitute a tightly integrated scene that can be manipulated as a unit. More generally, every space on the blend structure should have integration. In other words, the integration principle dictates that the blend must be recognized as a whole and as a new concept that is coherent.

OP2 *Intensifying Vital Relations*: compress what is diffuse by scaling a single vital conceptual relation or transforming vital conceptual relations into others.

OP3 *Maximizing Vital Relations*: create human scale in the blend by maximizing vital relations.

OP4 *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. Put differently, the topology principle dictates that every element projected into the blend should maintain the same neighborhood relations as in the input space. This principle can be disregarded without having a major impact in the value of the blend, especially if we are dealing with free combinations, such as an imaginary object with a given goal (Pereira 2005).

OP5 *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.

OP6 *Unpacking*: the blend alone must enable the blend reader/observer to unpack the blend to reconstruct the inputs, the cross-space mapping, the generic space, and the network of connections between all these spaces. Unlike other principles, unpacking takes the perspective of the blend reader, i.e., someone who is not acquainted with the blend generation process.

OP7 *Relevance*: all things being equal, if an element appears in the blend, there will be pressure to find significance for this element. Significance will include relevant links to other spaces and relevant functions in running the blend. In short, the relevance principle requires the existence of a reason for the blend to occur.

Blends and creative thought

The theory built around conceptual blending inevitably deals with the phenomenon of creative thinking. The ability of producing new, surprising, and valuable ideas or artifacts comes frequently in advanced forms of conceptual blending (Turner 2014). Due to the "intuitive" nature of creative thinking, the construction of a comprehensive theory of such phenomenon is quite challenging. Conceptual blending theory, without being an exception, is sometimes vague and less prone to formalization when dealing with crucial aspects of creative thought. In particular, the framework does not explicitly deal with novelty and the optimality principles do not clearly dictate whether a blend is creative or not. However, it is a common assumption that novelty can result from the application of these principles.

Despite all these limitations, the conceptual blending framework provides not only a set of sound principles but also a consistent terminology that can be used in creativity modeling. This has been a major motivation to consider the design of artificial creative systems based on computational approaches to conceptual blending.

Looking for good blends

By understanding what humans perceive as a good blend, we hope to dissipate some of the vagueness surrounding the explanation of parts of the blending mechanism that are not fully described or even ambiguous. We are primarily interested in analyzing the relevance of the optimality principles, the selection of input spaces as well as the projection of elements. Our goal is not to establish solid rules to the blending process – as it would be incongruous with the theory – but to provide some hints to questions such as "How 'semantically far' should the input spaces be to produce a good blend?","Is there a correlation between the quality of blends and the number of elements for projection?", or "Are all the optimality principles required to produce good blends?". In the case of artificial creative systems, we also expect to find clues to questions such as "Are the typical relationships found in concept maps sufficient to infer the quality of a blend?", "How important is to include common sense knowledge (sensorial and subjective elements) in concept maps to achieve better blends?", or "To what extent is required to have a goal-driven blending to obtain better blends?".

It should be noted that we do not use any a priori definition of what is considered to be a good blend. Constructing such a definition is actually the goal of this study. Nonetheless, our work relies on the premise that good blends are creative to some extent, whereas the reciprocal is not necessarily true.

In this paper, we focus on visual blends. More accurately, we work with images depicting fictional hybrid animals. Examples of hybrid animals, such as *Pegasus* or *the lion man*, are often presented in the literature as well-known and/or ancient blends. There are also several experiments in the field of computational creativity involving the creation of hybrid animals (Pereira and Cardoso 2003; Neahus et al. 2014). In our case, we opted to analyze this particular type of blend due to the fact that hybrid animals tend to be easily perceived as a blend, i. e., the blend reader can recognize the input spaces and simultaneously identify a novel creature. Nevertheless, we will try to make generalizations from our observations rather than drawing conclusions that only hold for this type of blends.

The survey

To assess the quality of blends, we conducted an online questionnaire survey in which approximately 100 participants judged 15 novel animals which are the result of blending anatomies from two different animals.[1] Each hybrid animal was depicted in one image/scene (see Figure 2). The author of all images but two is Arne Olaf (`http://gyyp.imgur.com/`). He uses *Adobe*® *Photoshop*® to create the hybrid creatures. The input images are put in two layers, adjusted in terms of size and unnecessary regions are removed. After that, he applies some common image processing techniques to make the transitions smoother.

Note that our focus is on blending at the conceptual level, overlooking aspects related to technical perfection. However, we are aware that technical perfection of a picture plays an important role in the perception of visual blends. This is why we decided to use blends with a similar level of quality in this respect - all chosen blends could be perceived as "good" as far as visual presentation is concerned. Moreover, the pictures share a similar rendering style. This enabled us to investigate other influential factors with more certainty as we ruled out rendering or poor presentation as a reason for bad human perception of a blend. This is particularly important in looking for findings that would hold also for other types of blends, not just the visual ones.

One may comment that there are no obviously bad blends in the dataset. This decision was based on a preliminary test done by ourselves, in which we noticed that there were big individual differences in the acceptance of blends, although the blends were all looking "nice" and the dataset presented good candidates for being well accepted. Our voting on blends was almost never unanimous, and this is why we wanted to investigate more thoroughly what could be expected on a bigger and more heterogenous population of subjects. With this in mind, "the good" and "the bad" blends from the title should be understood as "well accepted" and "not so well accepted" blends, showing the way towards creation of blends that will be well accepted by humans.

The criteria used in the survey cover some of the optimality principles criteria (e.g., by asking about coherence and consistency we are checking if a blended creature is perceived as having its own identity and corresponds to the optimality criterion of integration) as well as some criteria that define creativity, i.e., novelty, surprise, and value (Boden 1991). Thus, for each image depicting a hybrid animal, we asked the participants to rate the following criteria in a integer scale from 1 (the worst) to 5 (the best):

OI *Overall impression*;

N/S *Novelty/Surprise*;

I *Interestingness*;

AA *Aesthetic appeal*;

C/H *Comicality/Humor*;

C/C *Coherence/Consistency*;

PF *Evoques positive feelings*;

NF *Evoques negative feelings*;

CIP *Creative industries potential*.

The participants were also asked to provide a name to the hybrid creature as well us to inform us if they could easily recognize two distinct animals in the image. The latter question evokes the unpacking principle, i.e., the ability of the participant to reconstruct the input spaces.

Survey results and discussion

Figure 3 depicts the median overall impression for each of the hybrid animals in the dataset. According to these results,

[1] Available at `http://animals.janez.me`.

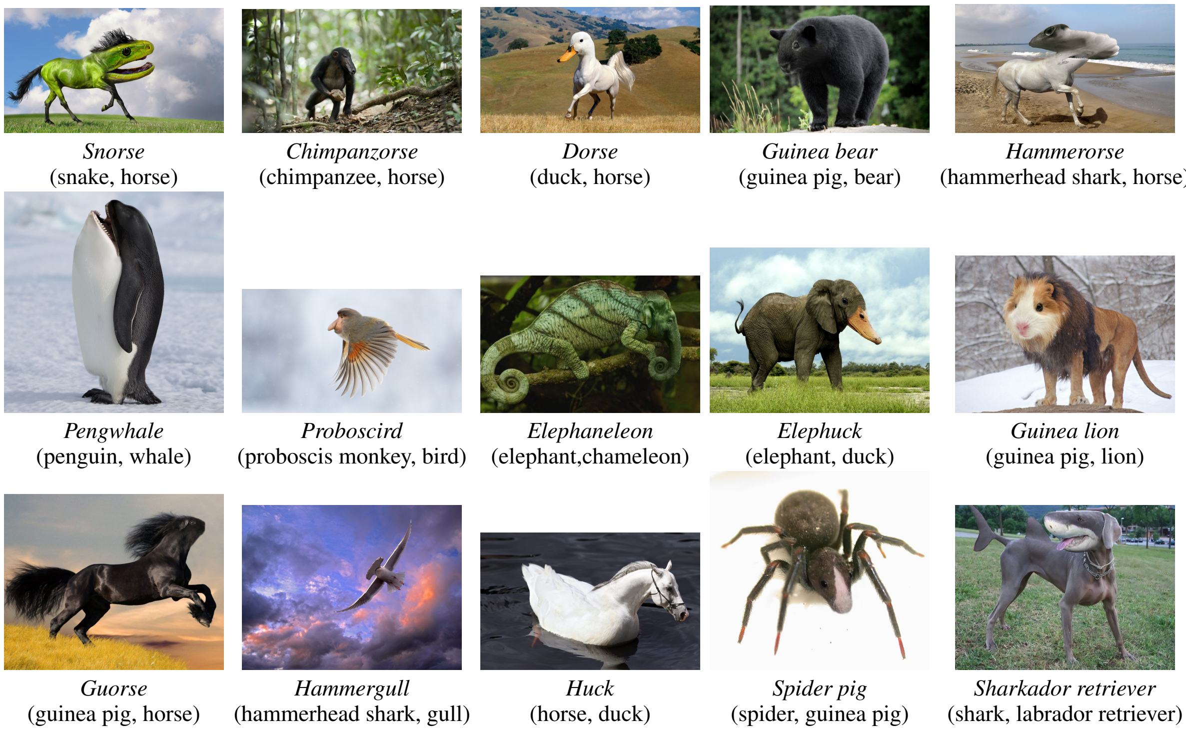

Snorse (snake, horse) · *Chimpanzorse* (chimpanzee, horse) · *Dorse* (duck, horse) · *Guinea bear* (guinea pig, bear) · *Hammerorse* (hammerhead shark, horse)

Pengwhale (penguin, whale) · *Proboscird* (proboscis monkey, bird) · *Elephaneleon* (elephant,chameleon) · *Elephuck* (elephant, duck) · *Guinea lion* (guinea pig, lion)

Guorse (guinea pig, horse) · *Hammergull* (hammerhead shark, gull) · *Huck* (horse, duck) · *Spider pig* (spider, guinea pig) · *Sharkador retriever* (shark, labrador retriever)

Figure 2: Hybrid animals dataset used in the online questionnaire. Each sub-caption contains the corresponding name of the blend as well as the input spaces. Names were coined by the authors of this paper or by the authors of the images and were not visible to survey participants. All blends were created by Arne Olaf, with the exception of *Sharkador retriever* and *Elephaneleon*, whose authorship is unknown. For a better visualization, some images were slightly cropped.

the top six best blends are *Guinea lion*, *Pengwhale*, *Guinea bear*, *Elephaneleon*, *Proboscird*, and *Dorse*, while *Spider pig*, *Hammerorse*, and *Guorse* were the least favorite blends.

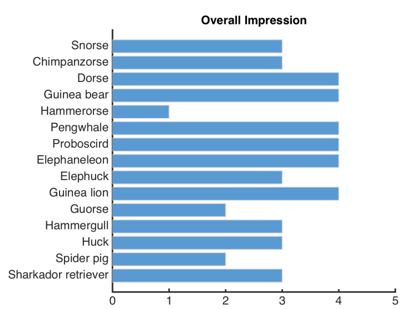

Figure 3: Overall impression (median) for each hybrid animal.

Figure 4 depicts a more detailed central tendency analysis of the survey results by including the median score for each criterion. Among the six best blends in terms of overall impression, *Guinea lion* and *Pengwhale* achieve the best overall scores. Out of the six best blends, five could be characterized by having a relatively big difference in the size of the original animal. It is also worth mentioning that *Guinea lion*, *Elephaneleon*, and *Pengwhale* are all in the best group with regard to the following criteria: novelty/surprise, interestingness, and coherence/consistency.

Regarding novelty, most blends achieved a median score of 4. The exceptions are *Chimpanzorse*, *Guinea bear* , *Elephuck*, *Guorse*, *Hammergull*, and *Huck*, which achieved a median score of 3. As it can be observed, high novelty does not necessarily lead to high overall impression. For instance, *Guinea bear* is a top-rated blend in terms of overall impression; however, its score in terms of novelty is among the lowest in the whole group. Conversely, *Hammerorse* has a high novelty score but a low overall impression score.

As pointed out by some respondents, novelty became more difficult to judge after a few images, as there were similar blends either in terms of input spaces or in terms of the elements for projection. This repetition and the fact that images were shown in fixed order might partially explain the lower scores obtained by *Elephuck*, *Hammergull*, or *Huck*.

Blends with a high overall impression score tend to have a high interestingness score. In fact, among the animals with the highest overall impression scores, *Dorse* is the only animal with an interestingness score of 3.

As for aesthetic appeal, we observe that blends with a low aesthetic appeal have a low overall impression score, whereas the most aesthetically appealing ones tend to have high overall impression scores.

Coherence/consistency scores tend to be well aligned with the overall impression. The animals with the lowest overall impression scores – *Hammerorse*, *Guorse*, and *Spider pig* – have a consistency score of 2. For the remaining blends, with the exception of *Dorse*, the median coherence score coincides with the median overall impression score. We also observe that animals with higher overall impression scores tend to evoke more positive feelings, while animals with lower overall impression scores tend to evoke more negative sentiment.

Creative industries potential scores are not always in concordance with overall impression results. Similar results can be observed for the criterion comicality/humor. However, blends with the lowest overall impression scores are seen as having a low creative industries potential.

These results clearly show that the novelty alone does not guarantee the overall rating nor creative industry potential nor how interesting the blends are. The one considered to be one of the most novel ones is *Hammerorse*, but its has the lowest overall rating of all. Similarly, *Smorse* and *Sharkador retriever* are among the most novel ones; however, this is not reflected in their overall impression scores.

Another statistic analysis is given in Figure 5, which contains the correlation among pairs of criteria. As it can be readily seen, aesthetic appeal is strongly correlated with the overall impression (ρ= 0.8). There is also a strong positive association between overall impression and coherence (ρ=0.76). This result reflects the importance of the optimality principles, as they are responsible for defining coherent blends. The correlation between novelty and overall impression (ρ=0.47) corroborates our previous remarks: it is difficult to establish a straightforward association between these two scores.

We received also more than 20 comments related to the questionnaire. The majority of them were expressing satisfaction (having fun, enjoying the survey, etc.). The negative comments related especially to the fact that the survey was too long and that it got monotonous after a while. Specific points were commented, such as that coherence was difficult to judge and that novelty and humor were not applicable after the first few images. It was also proposed that comparing more animals at the same time would be better. A few people also explained which were their favorites, with *Pengwhale* being mentioned a few times. Some people provided more original explanations, e.g., "The horse duck was boring, because they are both vegetarian".

The interview

We also conducted an interview with 4 people who took part in the survey. The main goal of the interview was to try to understand and discuss some of the ratings given by these participants. There was a general consensus that aesthetic appeal was an important requirement. For example, blending animals with similar types of coat – in terms of color, texture, or pattern – tends to result in aesthetically appealing blends. *Guinea bear* and *Guinea lion* were given as an examples of aesthetically appealing blends. *Snorse* was mentioned by one of the interviewees as another example of an aesthetically appealing blend, as there were no major differences between the snakeskin in the "snake part" and the coat in the "horse part" of the animal. *Pengwhale* was also a favorite among these participants. They enjoyed the fact that it was very difficult to establish a clear separation between "the whale part" and the "penguin part".

Participants took into account proportions when evaluating the aesthetic appeal. They presented *Guorse* as an example of a badly-proportioned blend: the proportions in the body of the horse require a head more elongated than the one of a guinea pig. In *Hammerorse*, the participants observed another instance of badly-proportioned parts. In this case, the head was seen as being too wide for the rest of the body.

One of the interviewees said to prefer the blend *Dorse* over the creature *Huck* because the head of *Dorse* has more resemblance with the head of a horse than the head of *Huck* has with the head of a duck. The interviewees also shared the opinion that surprise was required, but only to a certain extent. *Hammerorse* and *Spider pig* were given as examples of "too much surprise", which has a negative effect on the overall evaluation, whereas *Guinea bear* was presented as a blend with a minimal level of surprise.

Some participants suggested *Guinea lion* as a good example of comicality/humor due to the contrasting personalities of the animals given as input spaces. Although these mental spaces correspond to animals with similar coat and not so different anatomies, one is seen as a fierce creature, while the other one is a small harmless rodent

The participants emphasized the importance of recognizing the input spaces. However, there was the general idea that they enjoyed more when unpacking took time to occur.

Good blends: input spaces, projection, and optimality principles

The level of novelty or surprise in a blend is partially dictated by the selection of input spaces and the choice of elements for projection. While the results from the survey do not show a direct association between novelty/surprise and the overall impression of the blend, it is somehow clear that both novelty and surprise are required to some extent. Selecting seemingly unrelated input spaces seems a good option only if the choice of elements for projection and subsequent tasks are able to deconstruct the idea that both concepts are unrelated. In this particular case, projection should be able to highlight various links between the two mental spaces that are less obvious instead of establishing a reduced number of more obvious connections.

Figure 6 depicts the concept similarity between different concepts used in the survey. Instead of using Linnaean taxonomy to compute the similarity between two animals, we opted for a more generic and elaborate measure that is able to generate more fine-grained results.[2] The concept similarity was calculated by applying the Personalized PageRank (PPR) (Haveliwala 2003) to ConceptNet. PPR is a variation of the standard PageRank algorithm used to rank nodes in a

[2] In our experiments with Linnaean taxonomy, the distances between animals were 5, 6, or 7.

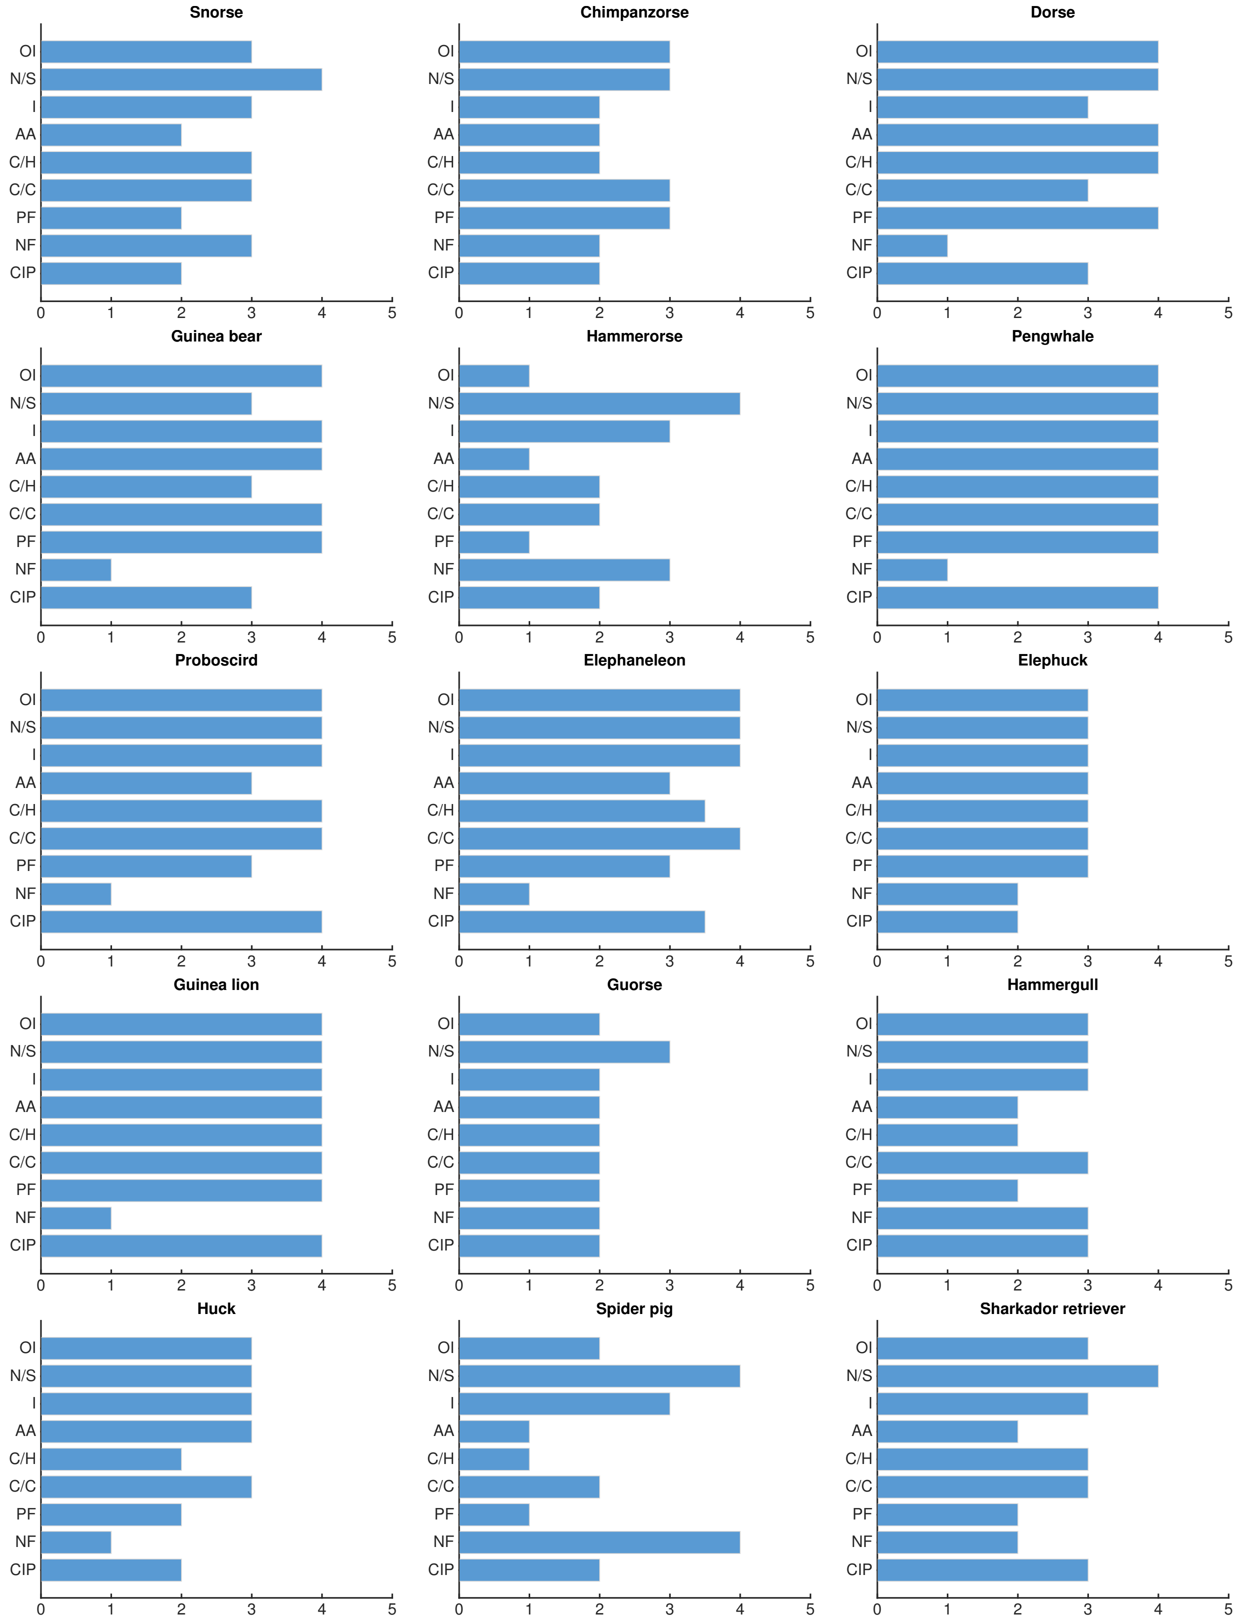

Figure 4: Survey results for each one the 15 hybrid animals. The bars represent the median score for each of the criteria.

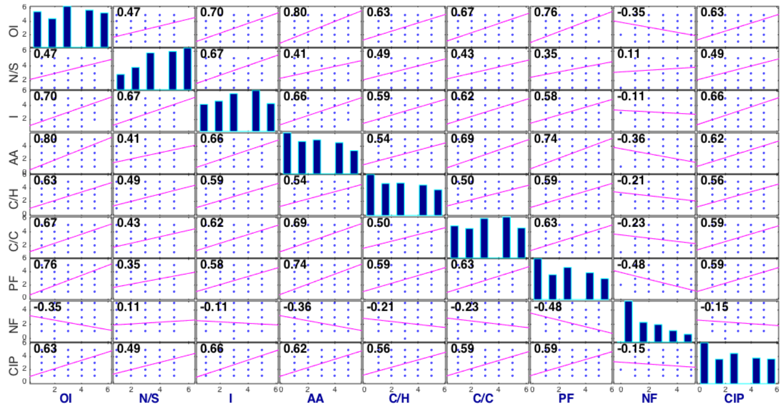

Figure 5: Matrix depicting correlations among pairs of criteria used in the survey. Non-diagonal elements contain scatter plots of the variable pairs. Diagonal elements contain histograms of the variables. The slopes of the least-squares lines in the scatter plots correspond to the displayed correlation coefficients.

network (Page et al. 1999). The PPR of a node v in a network (PPR_v) is a vector which, for each other node w in the network, tells how simple it is to randomly walk from v to w. It is calculated as a stationary distribution of the position of a random walker which starts its walk on node v and at each step either (with probability p) randomly selects one of the connections leading out of its current node and travels along it or (with probability $1-p$) travels back to its starting location. In our experiment, p was set to 0.85.

If the PPR of node w according to node v ($PPR\_v(w)$) is high, this means that the node w is easy to reach from the departing node v. However, the path from v to w is not symmetric to the one from w to v. Therefore, the similarity measure s is proposed, where $s(v,w) = PPR\_v(w) + PPR\_w(v)$. In short, the higher the score the stronger the connection between the nodes and the higher the similarity between concepts.

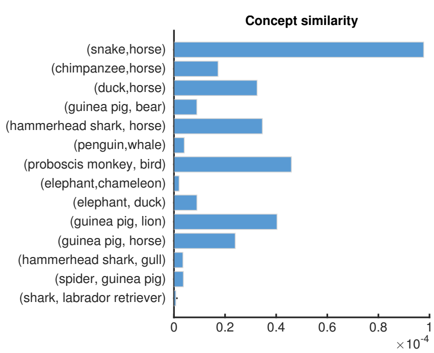

Figure 6: Similarity between the input spaces used in the survey.

For our work, we used the ConceptNet graph to calculate the similarity between two animals. We ran the PPR algorithm on the network to obtain the personalized PageRank vector for each of the animals in question. The personalized PageRank of a vertex is calculated iteratively by spreading the rank of the original vertex along its connections until the rank is no longer substantially changing.

This metric cannot be straightforwardly associated with the overall impression or novelty scores, as it does not faithfully reflect how semantically far the concepts are for a given observer. However, it suggests that sometimes seemingly unrelated input spaces (e.g, a horse (mammal) and a snake (reptile)) are sometimes more similar than two mammals (e.g., guinea pig and bear). We believe that exploring these less obvious similarities is a good starting point for the construction of high-quality blends.

Not unexpectedly, the results from the survey support the idea that all optimality principles are relevant, with the exception of topology (as already explained in the previous section). Integration is arguably the most important one and it should not be overlooked. It is necessary (but not sufficient) to dictate the coherence of the blend.

Figure 7 shows the percentage of affirmative answers to the unpacking question: "Can you easily recognize two distinct animals in the image?" for each one of the blends. In general terms, input spaces were easily recognized, although this task became more difficult when unpacking *Guinea bear* and *Guorse*, as the differences between the animal that provides the body and the blend are minimal. We believe that unpacking is a relevant principle, but it should not be given priority over other principles such as integration. On one hand, it allows the blend reader/observer to build his own interpretation of the blending process, which is fundamental to preform assessments from the perspective of the reader/observer. On the other hand, an immediate unpacking sometimes means a lack of surprise or novelty.

Conclusions and Future Work

We presented and discussed an evaluation based on the human perception of visual blends. This research is part of a wider study which is oriented towards two major aims: (i)

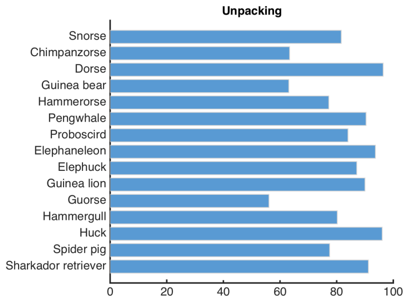

Figure 7: Unpacking: percentage of participants who responded affirmatively to the question "Can you easily recognize two distinct animals in the image?".

to help clarify some less clear and less described aspects of the blending mechanism which play a fundamental role in creative thinking; (ii) to improve the creation and curation processes in artificial creative systems.

Although we have only dealt with visual blends depicting hybrid animals, some of our observations can be applied to other types of blends. For instance, surprise and novelty are necessary but not sufficient to guarantee a high-quality blend. In fact, too much surprise is unfavorable if it affects the consistency of the emerging structure. The survey results also reflect the importance of having coherent blends, which emphasizes the importance of the optimality principles.

In this first experiment, we inevitably dealt with the specificities of visual blends, all being of similar technical quality, depicting hybrid animals. The results demonstrated that aesthetic appeal is an important criterion. Besides the quality of rendering, there are other aspects, namely symmetry and proportions, that influence aesthetic appeal. This may not be a relevant criterion when analyzing non-visual blends. However, since aesthetic appeal is related to symmetry and proportions, we argue that this criterion should be considered even when we are not working in the visual domain. For this reason choosing blends of similar technical quality, even at the cost of lower variety on the scale of all possible blends, seems to be the right decision, if we want to gain more insight into the conceptual level of blending. An interesting question remaining for future work is whether the results would be different if only textual descriptions or concept maps were given to the test subjects.

While the correlation of overall scores with other criteria in our experiment helps to identify the blends perceived as good or bad, the AHA! effect is correlated to the level of novelty, surprise, unpacking and creative industry potential. This will be further investigated with the analysis of the names given to the blends by the test subjects. Some of these names were very creative and reflected new qualities, existing in the blend while not being present in the input spaces. This will help us to understand the role of the emergent new structure reflected in such names, and might uncover the potential of blends to trigger the highly individual AHA! effect and human creativity.

Acknowledgements

The authors acknowledge the financial support from the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under the ConCreTe FET-Open project (grant number 611733) and the PROSECCO FET-Proactive project (grant number 600653). The authors thank also Jan Kralj for designing the concept similarity measure, and Janez Kranjc for the survey platform, which was developed within the WHIM FET-Open project (grant number 611560). We would also like to thank the survey participants for their time and their valuable comments.

References

Boden, M. A. 1991. *The Creative Mind: Myths and Mechanisms*. Basic Books, Inc.

Colton, S.; Pease, A.; and Ritchie, G. 2001. The effect of input knowledge on creativity. In *Proceedings of the First Workshop on Creative Systems, International Conference of Case-Based Reasoning (ICCBR'01)*.

Colton, S. 2008. Creativity versus the perception of creativity in computational systems. In *Proceedings of the AAAI Spring Symp. on Creative Intelligent Systems*.

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.

Haveliwala, T. H. 2003. Topic-sensitive pagerank: A context-sensitive ranking algorithm for web search. Technical Report 2003-29, Stanford InfoLab.

Jordanous, A. 2012. *Evaluating computational creativity: a standardised procedure for evaluating creative systems and its application*. Ph.D. Dissertation, University of Sussex.

Joy, A.; F. Sherry Jr., J.; and Deschenes, J. 2009. Conceptual blending in advertising. *Journal of Business Research* 62(1):39 – 49.

Neahus, F.; Kutz, O.; Codescu, M.; and Mossakowski, T. 2014. Fabricating monsters is hard - towards the automation of conceptual blending. In *Proceedings of the Workshop "Computational Creativity, Concept Invention, and General Intelligence"*.

Page, L.; Brin, S.; Motwani, R.; and Winograd, T. 1999. The pagerank citation ranking: Bringing order to the web. Technical Report 1999-66, Stanford InfoLab.

Pereira, F., and Cardoso, A. 2003. The horse-bird creature generation experiment. *The Interdisciplinary Journal of Artificial Intelligence and the Simulation of Behaviour(AISBJ)* 1(3):257–280.

Pereira, F. 2005. *Creativity and AI: A Conceptual Blending approach*. Ph.D. Dissertation, University of Coimbra.

Ritchie, G. 2001. Assessing creativity. In *Proceedings of the AISB'01 Symposium on Artificial Intelligence and Creativity in Arts and Science*, 3–11.

Turner, M. 2014. *The Origin of Ideas*. Oxford University Press.

Wiggins, G. 2001. Categorising creative systems. In *Proceedings of the Third Workshop on Creative Systems (IJCAI'03)*.