Deciphering The Cookie Monster: A case study in impossible combinations

Maria M. Hedblom
Institute of Artificial Intelligence
University of Bremen, Germany
hedblom@uni-bremen.de

Guendalina Righetti
KRDB, Computer Science Faculty
Free University of Bozen-Bolzano, Italy
guendalina.Righetti@stud-inf.unibz.it

Oliver Kutz
KRDB, Computer Science Faculty
Free University of Bozen-Bolzano, Italy
okutz@unibz.it

Abstract

In conceptual blending, the transfer of properties from the input spaces relies on a shared semantic base. At the same time, interesting blends are supposed to resolve deep semantic clashes where many concept combinations correspond to impossible blends, i.e. blends whose input spaces lack any obvious similarities. Instead of a shared structure, the blends are based on bidirectional affordance structures. While humans can easily map this information, computational systems for creative constructions require an understanding of how these features relate to one another. In this paper, we discuss this problem from the perspective of linguistics and computational blending and propose a method combining theory weakening and semantic prioritisation. To demonstrate the problem space, we look at the Sesame Street character 'The Cookie Monster' and formalise the blending process using description logic.

Introduction

Conceptual blending (CB) has been proposed as a model of combinational creativity (Boden, 1998; Fauconnier and Turner, 2008). Based on the principle of analogical transfer, where information of one domain is transferred onto another based on their shared structure, CB suggests that creativity arises as a conceptual merge of two input spaces. Cognitively speaking, blending is a dynamic process guided by a series of optimality principles that repeatedly update the interpretation of both the shared structure and the blended concept (Fauconnier and Turner, 2008). In computational approaches to CB, a static knowledge representation is required, limiting the possibilities to model the dynamics of emergent processes in cognition (e.g. Kutz et al. (2014); Pereira and Cardoso (2002)). This becomes a problem when concepts from (very) different ontological branches are merged, a phenomenon called *impossible blends* (Turner, 1996). While humans can through mental elaboration principles find shared structure and connections between concepts that have little to nothing in common, computers are forced to rely on the information they are presented with and innovative methods to identify potentially shared structure are required in order to advance computational CB.

In comparison to computational blending, the creative process of noun-noun (NN) compound word constructions demonstrate how humans easily can blend words from different ontological branches based on other criteria than directly shared structure. Concepts like *coffee cup* and *face-palm* do not result from some manipulation of the intersection of their respective input spaces. Instead, a coffee cup is a cup particularly designated to contain coffee based on the inverse role and bidirectional affordances of how a cup is a container for liquids and how liquids need a container. For a face-palm, there exists a shared ontological structure, namely body-part. However, a face-palm is not a body-part, but rather slang for the emotional reaction captured in the embodied action of covering your eyes with your hand. Computational blending systems that rely on mapping the intersection of two input spaces would not be able to reach these interpretations as they involve an understanding of embodied experiences that is not shared between the input spaces.

In order to formally deal with impossible blends, they need to be treated with respect to the semantic components within the respective input spaces and how they *could* relate to one another. As in the examples above, some important semantic components are object affordances (Gibson, 1977) and semantic components describing perceptive and embodied experiences, such as those found in image schemas (Johnson, 1987).

To learn how to better deal with impossible blends in computational blending, we present a top-down analysis of the creative process that takes place when conceptualising the impossible blend found in the Sesame Street character *The Cookie Monster*. Our method utilises ontological weakening of the input spaces to identify shared semantic structure with an emphasise on identifying transitive and inverse roles of affordances and by performing property interpretation in the form of semantic prioritisation. In a miniature setting, we formalise the spaces using Description Logic (DL).

All that combines is not Blending

Compound words are lexical compositions in which one domain inherits properties from another domain by merging two words, e.g. blackbird and coffee cup. In English, two parts can be distinguished: the Head, denoting the class, and the Modifier, restricting the meaning of the word. For instance, a compound concept such as Cookie Monster would be differently interpreted than a Monster Cookie:

Proceedings of the 12th International Conference on Computational Creativity (ICCC '21)
ISBN: 978-989-54160-3-5

in the first case the word Monster plays the role of the Head, modified by the word Cookie (possibly a *Cookie eating Monster*, see below). In the second one, it is the other way around, and the concept is more likely to be interpreted as something like a *Cookie* which is *monstrous* in some respects.

According to Wisniewski (1997), NN combinations can be interpreted in three ways: 1) The first is the *relation-linking* interpretation, where some kind of relation between the components is highlighted (the Cookie Monster is a monster that *eats* cookies). 2) The second is the *property* interpretation, where one or more properties of the Modifier noun apply to the Head concept (the Cookie Monster is a monster that is as *sweet* as cookies). 3) The third is called *hybridisation*, which is a "combination of the two constituents [...] or a conjunction of the constituents" (Wisniewski, 1997, p.169). The result of the combination corresponds essentially to a 'mash-up' or 'blend' of both components (the Cookie Monster would then be both a cookie and a monster).

However, even with this differentiation, the inheritance relationship from the input spaces is not always straightforward. Consider the difference between the NN compound words *snowman* and the *ice-cream man*. Ontologically, the input spaces snow and ice-cream share several properties such as being cold and fluffy, yet the compound words are ontologically distinct based on essential properties and what they are 'used' for. In the case of a *snowman*, the Modifier's properties are transferred in its entirety as the compound refers to *a man made out of snow*[1]. Hence, *snowman* corresponds to a hybridisation of the two concepts. In the *ice-cream man*, the result has little to do with any properties of ice-cream. Instead, the *ice-cream man* blend calls for a relation-linking interpretation on weakened input spaces based on functionality. Here, man: *ability to bring* is treated in relation to the ice-cream space, essentially making the blend *a man who brings ice-cream*.

In comparison to the linguistic research on compound words, CB is the emergent process that finds this intersection during (primarily) hybridisation. The blend inherits properties from both input spaces and through emergent properties and optimality principles, the blends are ensured to make sense from a cognitive perspective (Fauconnier and Turner, 1998; Pereira and Cardoso, 2003).

Turning such cognitive processes into 'artificially intelligent' identification of shared structure and projection of relevant information is a non-trivial problem. One important feature is that (most often) the most salient and semantically rich features should be inherited by the blend. Arguably, the essence of objects are tightly connected to the affordances they offer (Gibson, 1977), especially in terms of their functional, spatiotemporal behaviours.

Affordances have an interesting feature. They are essentially bi-directional dispositions (Beßler et al., 2020) with transitive or inverse roles of the participants (e.g. inverse roles in DL (Horrocks and Sattler, 1999)). For instance, Food has properties that offer the affordance *ToBeEaten* as its most essential property is to be edible. Simultaneously, a LivingCreature has the behaviour *CanEat* as it is essential it should eat, else it is not alive (for long). These kinds of essential properties are of crucial importance when performing CB and interpreting compound words and should, therefore, be incorporated in the formal blending process.

Two important suggestions for improving semantics for computational CB have been introduced (e.g. see (Eppe et al., 2018; Hedblom, 2020)). The first is *theory weakening,* in which the input spaces are ontologically generalised into spaces of less detail, or higher-order components, to better identify potentially shared structure. Axiomatised theory weakening has been applied in logical approaches to analogy and CB (e.g. Gentner (1983); Schmidt et al. (2014)). However, they lack semantic selection. The second suggestion is *semantic prioritisation*[2] that promotes that the most important attributes and properties should be transferred into the blend. An example is the property interpretation found in the *houseboat* blend: despite being a house to live in, it also moves on water as the most salient image-schematic affordances of the Modifier, boat, is projected into the blend. In the next section, we demonstrate how these two methods are involved in deconstructing the Cookie Monster blend.

The Complexity of the Cookie Monster

Actually named Sidney Monster, *The Cookie Monster* is a blue hand-puppet from The Muppets famous for his obsession with eating cookies. The conceptual complexity that emerges when looking at the Cookie Monster as a compound blend is the following: While the character is a Monster by Muppet-classification, the most appropriate interpretation is that the epithet *monster* is there to describe an unnatural relationship to cookies. Compare it with calling a (non-monster) friend a "cookie monster" if s/he eats a lot of cookies.

From a formal blending perspective, two interesting things happen in this blend:
1) The blended space is **not** an intersection between the input spaces Monster and Cookie. Instead, it corresponds to a relation-linking interpretation, based on a conceptual mapping between the inverse roles of: *edible* and *canEat*.
2) The second thing that happens is that the blend is not exclusively *a cookie-eating monster*. Cookie Monster is a sweet character that simply eats cookies in a *monstrous* way (over-consumption, guzzling, etc.). This is a form of property interpretation where the *sweetness* of the cookies are transferred onto the blend and, even more interestingly, the *unnaturalness* of monsters are transposed onto the *cookie-eating* property.

As cookies and monsters have nothing in common in their conceptual spaces, they need to be generalised to the point in which a connection can be made. For this, we use the

[1]Arguably it would be possible to claim that a snowman is just snow that inherited the shape of a man. However, we pertain that the snowman is inadvertently also given individual identity through anthropomorphism.

[2]Hedblom (2020) calls this *image schema prioritisation* and focus on spatiotemporal relationships. However, the idea can be transposed onto any conceptual components of semantic importance.

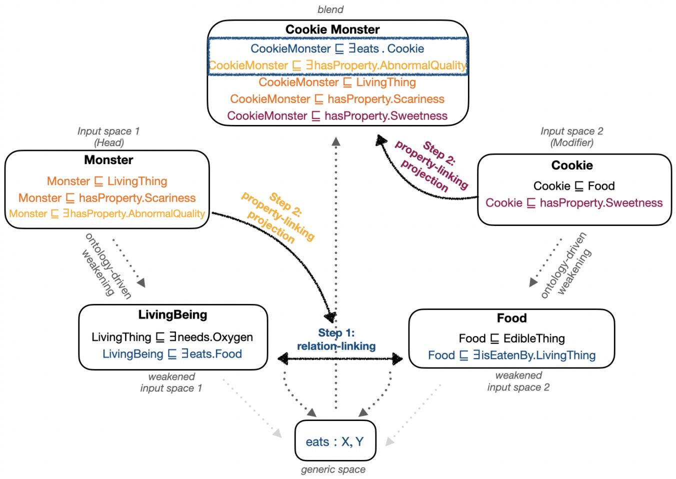

Figure 1: Blending diagram of The Cookie Monster

ontological branches to identify the first bi-directional affordance relationships and 'step up' in the object classes.

In Figure 1, the blending diagram for Cookie Monster is presented with respect to both theory weakening and semantic prioritisation based on property interpretation, explained in more detail below.

Initiation of Computing the Impossible

Step 1: Relation Linking through Weakening. A computational approach looking for a relation-linking interpretation needs to identify a relation that holds between the two input spaces. In our top-down example, and following a logic-based representation, this corresponds to looking for a role R holding between the instances of the classes in the input spaces. If such a role is identified, then the task would be easily solved. For instance, for the input spaces 'Human' and 'Monster' one trivial relation-linking interpretation would be R:scares as monsters are commonly perceived as dangerous to humans[3]. However, for the input spaces Monster and Cookie, ontologically represented in Figure 1, no such obvious relation exists[4].

Addressing this, one (or both) input spaces need to be generalised until a *shared role* is found. This corresponds to a form of theory weakening, where the *weakening* allows to 'step up' in the ontological branches by exploiting the subsumption relations holding between the concepts in the ontologies. One possibility to formally capture this is to utilise a generalisation operator as described in (Confalonieri et al., 2020) and as exploited in (Confalonieri and Kutz, 2020). In short, a generalisation operator with respect to an ontology is a function γ_O that takes a concept C and returns the set $\gamma_O(C)$ of the *super-concepts* of C[5]. Intuitively, a concept D is a generalised super-concept of the concept C with respect to an ontology O if in every model of the ontology all instances of C are also instances of D.

In our example, the ontological assumption presented in Figure 1 claims that LivingBeing is a super-concept of the concept Monster. Applying theory weakening, the input space Monster is then generalised into LivingBeing, and a relation holding between LivingBeing and Cookie is sought. If a relation is found, it is returned. Otherwise, as in this case, also the other input space needs to be generalised from Cookie into Food. Here, the role eats holds between the instances of LivingBeing and Food and constitutes the shared structure that belongs to the generic space.

Following blending heuristics, the information in the generic space constitutes the foundation for the blend by adding the specific information from the input spaces, generating the blend CookieMonster$\sqsubseteq$ $\exists$eats.Cookie. However, as CookieMonster$\sqsubseteq$Monster is also a correct interpretation, and following the transfer of information between Head and Modifier, the blend will also be defined by the axioms describing the Head ontology of Monster. Yet, one more complexity arises due to the nature of this impossible blend. This leads us to step 2.

Step 2: Semantic Property Prioritisation. Semantic prioritisation is a form of formal property interpretation. It suggests that the most salient features should be identified in the input spaces and inherited into the blend. One such interesting semantic transfer is the mapping of the abnormality of Monster into the role identified in the previous step, namely the eats function. This role is enhanced by an *abnormal* relationship, i.e. to define a role inclusion such as abnormallyEats $\sqsubseteq$ eats (Horrocks and Sattler, 1999).

Finally, as with all blending, the Modifier concept is there to alter the nature of the Head concept. For Cookie, the most salient feature, i.e. that distinguishing it from the other foods, is $\exists$hasProperty.Sweetness and its conceptual space extends more than just sugary foods—but also sweet and desirable characteristics. In contrast, monsters are *scary* and its conceptual space is directly inconsistent with that of sweetness. Directly transferring the salient features of both input spaces, therefore, may create a logical impossibility. Working top-down we already know that Cookie Monster is a charming fellow, hence priority is given to the Modifier. In unknown combinations or situations with multiple salient features, prioritisation strategies need to be applied to identify the most appropriate mapping.

Identifying the salient features to be transferred is one of the biggest challenges for future work.

Discussion and future work

Computational blending has become one of the most widely used methods for simulating computational creativity, yet human ability still far exceeds the current state of the art of computational systems.

To contribute to this research agenda, we took a brief look into impossible blends by using the Cookie Monster

[3]Based on the assumption that monsters are inherently scary (Neuhaus et al., 2014).

[4]In the context of a formal ontology Monster$\sqsubseteq$ LivingBeing and LivingBeing$\sqsubseteq$ $\exists$ eats.Food implies Monster$\sqsubseteq$ $\exists$ eats.Food. Ideally, theory weakening could be set to exploit logical inference directly.

[5]Conversely, also a *specialisation operator* can be defined.

as a case study of a compound of two ontologically distinct branches. Building on linguistic research on noun-noun compound words, one of the paper's main contributions to formal CB is showing how theory weakening could be employed to identify relation-linking interpretations, as well as utilising semantic prioritisation of salient features to more accurately deal with property projection.

The ideas follow the large body of work aiming to improve computational blending (e.g. Eppe et al. (2018); Neuhaus et al. (2014); Veale, Seco, and Hayes (2004)). The distinction between Head and Modifier is also reminiscent of asymmetric amalgams as described in (Besold, Kühnberger, and Plaza, 2017), with the difference that our approach does not identify a traditional generic space and uses different computational techniques as well.

Many challenges remain in order to utilise these ideas to deal with impossible blends in computational CB. To address these, future work includes incorporating the work on generalisation operators together with a system for semantic prioritisation. More precisely, by building on previous research and on empirical results regarding concept salience in compound words (Devereux and Costello, 2012), we plan to combine the formal work on inverse and transitive roles (Horrocks and Sattler, 1999) together with an ontological repository of affordances and other semantic components.

Acknowledgements

Special thanks to Prof. Mark Turner for providing valuable input on conceptual blending during the completion of this research.

The research reported in this paper has been supported by the German Research Foundation DFG, as part of Collaborative Research Center (Sonderforschungsbereich) 1320 "EASE - Everyday Activity Science and Engineering", University of Bremen (http://www.ease-crc.org/).

References

Besold, T. R.; Kühnberger, K.-U.; and Plaza, E. 2017. Towards a computational-and algorithmic-level account of concept blending using analogies and amalgams. *Connection Science* 29(4):387–413.

Beßler, D.; Porzel, R.; Pomarlan, M.; Beetz, M.; Malaka, R.; and Bateman, J. 2020. A formal model of affordances for flexible robotic task execution. In *Proc. of the 24th ECAI*.

Boden, M. A. 1998. Creativity and artificial intelligence. *Artificial intelligence* 103(1-2):347–356.

Confalonieri, R., and Kutz, O. 2020. Blending under deconstruction. *Annals of Mathematics and Artificial Intelligence* 88(5):479–516.

Confalonieri, R.; Galliani, P.; Kutz, O.; Porello, D.; Righetti, G.; and Troquard, N. 2020. Towards even more irresistible axiom weakening. In Borgwardt, S., and Meyer, T., eds., *Proc. of DL*, volume 2663. CEUR-WS.org.

Devereux, B. J., and Costello, F. J. 2012. Learning to interpret novel noun-noun compounds: Evidence from category learning experiments. In *Cognitive aspects of computational language acquisition*. Springer. 199–234.

Eppe, M.; Maclean, E.; Confalonieri, R.; Kutz, O.; Schorlemmer, M.; Plaza, E.; and Kühnberger, K.-U. 2018. A computational framework for conceptual blending. *Artificial Intelligence* 256:105–129.

Fauconnier, G., and Turner, M. 1998. Conceptual integration networks. *Cognitive Science* 22(2):133—187.

Fauconnier, G., and Turner, M. 2008. *The way we think: Conceptual blending and the mind's hidden complexities*. Basic Books.

Gentner, D. 1983. Structure mapping: A theoretical framework for analogy. *Cognitive Science* 7(2):155–170.

Gibson, J. J. 1977. The theory of affordances. In Shaw, R., and Bransford, J., eds., *Perceiving, Acting, and Knowing: Toward an Ecological Psychology*. Hillsdale: NJ: Lawrence Erlbaum. 67–82.

Hedblom, M. M. 2020. *Image Schemas and Concept Invention: Cognitive, Logical, and Linguistic Investigations*. Cognitive Technologies. Springer Computer Science.

Horrocks, I., and Sattler, U. 1999. A description logic with transitive and inverse roles and role hierarchies. *Journal of logic and computation* 9(3):385–410.

Johnson, M. 1987. *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*. Chicago and London: The University of Chicago Press.

Kutz, O.; Bateman, J.; Neuhaus, F.; Mossakowski, T.; and Bhatt, M. 2014. E pluribus unum: Formalisation, Use-Cases, and Computational Support for Conceptual Blending. In *Computational Creativity Research: Towards Creative Machines*, Thinking Machines. Atlantis/Springer.

Neuhaus, F.; Kutz, O.; Codescu, M.; and Mossakowski, T. 2014. Fabricating monsters is hard: towards the automation of conceptual blending. In *Proc. of C3GI*. Citeseer.

Pereira, F. C., and Cardoso, A. 2002. Conceptual Blending and the Quest for the Holy Creative Process. In *Proc. of the 2nd Workshop on Creative Systems*.

Pereira, F. C., and Cardoso, A. 2003. Optimality principles for conceptual blending: A first computational approach. *AISB Journal* 1(4):351–369.

Schmidt, M.; Krumnack, U.; Gust, H.; and Kühnberger, K.-U. 2014. Heuristic-Driven Theory Projection: An Overview. In Prade, H., and Richard, G., eds., *Computational Approaches to Analogical Reasoning: Current Trends*, volume 548 of *Computational Intelligence*. Springer-Verlag.

Turner, M. 1996. *The literary mind: The origins of thought and language*. Oxford University Press.

Veale, T.; Seco, N.; and Hayes, J. 2004. Creative discovery in lexical ontologies. In *Proc. of the 20th COLING Computational Linguistics*, 1333–1338.

Wisniewski, E. 1997. When concepts combine. *Psychonomic Bulletin & Review* 4:167–183.