Introducing Pathomalgametry:
Conceptual Blending with Geometric Path-Finding and Amalgamation

Mena Leemhuis
Free University of Bozen-Bolzano
Bozen/Bolzano, Italy
mena.leemhuis@unibz.it

Oliver Kutz
Free University of Bozen-Bolzano
Bozen/Bolzano, Italy
oliver.kutz@unibz.it

Abstract

Conceptual blending, where new concepts are created through a selective combination of known ideas, is a widely known option for concept invention within computational creativity research. However, existing approaches implementing blending are often either neglecting conceptual aspects, as, e.g., in image morphing, or they suffer from the high complexity of the creation of the blend. Therefore, we propose a new neuro-symbolic approach for conceptual blending of ontologies, based on knowledge-graph embeddings. Here, the inherent structure of the embedding space is used both to identify a generic space and to guide the blending process by interpreting blending as path search in the embedding space by iteratively relaxing the input concepts. Our approach is efficient as it can reuse existing embedding techniques as created for, e.g., link prediction. Its suitability is discussed both theoretically and based on a toy example.

Introduction

One well-known form of human creativity has been described by Fauconnier and Turner (2002) in the framework of *conceptual blending* (CB for short), namely the selective combination of known ideas into a ‘blend’, a surprising and creative new combination of selected features of the inputs. One well-known example is, e.g., the “houseboat” that can be seen as a conceptual blend of “house” and “boat”. Although Fauconnier and Turner introduced blending as a cognitive process, there has been significant effort over the past two decades to find computational realizations in order to facilitate computational creativity. Approaches are, e.g., interpreting it as a search problem solved with answer set programming (Eppe et al. 2018), a formalization based on algebraic specification (Goguen and Harrell 2010), or using mappings and structural alignments (Pereira 2007).

Here we focus on CB of concepts based on *amalgamation* (as, e.g., in (Confalonieri et al. 2018; Confalonieri and Kutz 2020)), thus by relaxing two input concepts successively until a common *generic space* is reached. The combination of two of these relaxations then results in a blend. However, symbolic approaches to conceptual blending in general suffer from high computational complexity of determining the generic space and the best suited blend (see, e.g., the discussion in Eppe et al. (2018)). As conceptual blending is considered a highly intuitive process (Fauconnier and Turner 2002), it seems to be natural to consider it not on a purely symbolic level but to incorporate subsymbolic information, e.g., in the form of similarity information. This has been done, e.g., by Wang et al. (2023), performing conceptual blending in large language models using their associative capacities. However, this approach does not focus on consistency and optimality principles in the rigidity needed for CB in an ontological context. Another approach was presented by Olearo et al. (2024) based on the idea that a blend is in between the latent space representations of the input concepts in a diffusion model. Though the approach incorporates the implicit information learned by the diffusion model and thus models some form of salience of the features, it neither considers a generic space nor conceptual information explicitly. Therefore, a different approach is needed, combining the following two ideas:

(A) the conceptual information should be respected to ensure consistency and coherence with the basic principles of conceptual blending, and

(B) subsymbolic information should inform the blending process by incorporating information regarding concept similarity.

An area well known for combining symbolic ontology information with subsymbolic similarity information is *knowledge base embedding* (see, e.g., (Bourgaux et al. 2024)). There, symbolic information is embedded into a geometric space by representing concepts as (typically) convex regions, instances as points and relations and logical operations as geometrical operations. One main use case is link prediction, thus the prediction of relations between instances. Another is concept membership prediction which allows to reduce the incompleteness of an ontology and to adapt to specific domains. An embedding can be used as a guide to determine a suitable generic space and consequently a suitable blend: As similar concepts are modeled closer to each other, the generic space can be interpreted geometrically as the smallest region containing both input concepts. This avoids over-generalization, and defines a region of possible blends. For determining the best out of these possible blends, the viewpoint of Leemhuis and Kutz (2024) can be applied to interpret the blending process as a path-finding problem in this space. The basic idea is that the generic space can be represented as the smallest convex set incor-

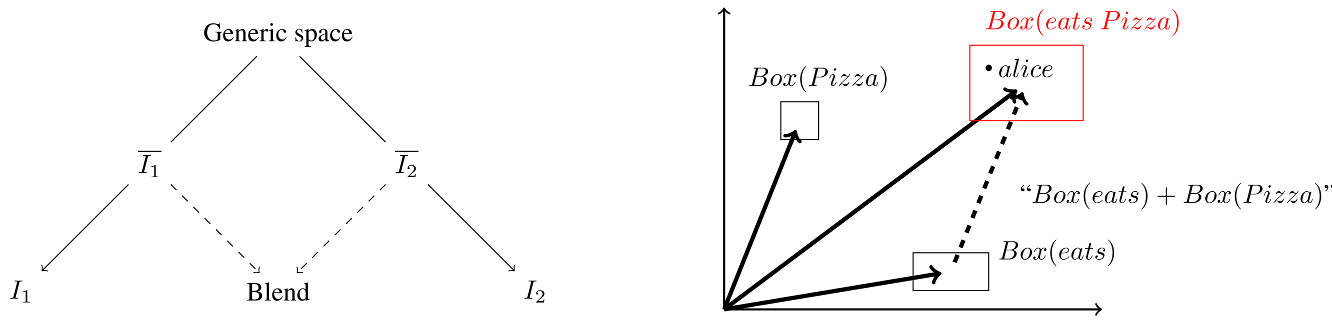

Figure 1: (a) basic diagram of CB with amalgams; (b) example for TransBox representing the fact that "Alice eats Pizza".

porating the representation of both input concepts. A blend can be considered as "in between" the input concepts on a path through the generic space and thus preserves some of the meaning of both input concepts. This eases the blending process and allows for incorporating optimality principles into the process implicitly. In the following, a first version of this neuro-symbolic conceptual blending approach is presented, and its suitability and features are discussed both theoretically and with the help of a toy example.

Preliminaries

Ontology-based Approaches for Blending

CB is a framework describing the creative invention of a new concept by combining known ones. The shared generalization of the known concepts, the *generic space* (GS), is found and is used as basis to define a blend, thus a combination of salient features of the input concepts, adhering to optimality principles (Fauconnier and Turner 2002). CB is applied in many different contexts, e.g., in poetry or even for mathematical concepts. Here, we are focusing on ontological blending, thus interpreting the inputs of the blending process as concepts of an ontology. Ontologies are used to represent structured information of the world in form of axioms representing general knowledge and assertions representing facts. One way of representing ontologies is with the help of description logics (DLs) (Baader et al. 2007). We are focusing here on ontologies in the DL $\mathcal{EL}_\bot$(a fragment of the well-known $\mathcal{EL}^{++}$(Baader, Brandt, and Lutz 2005)) due to its computational advantages, as subsumption is polynomial. It includes concept conjunction, existential quantification, subsumption, and especially disjointness of concepts.

One idea is to interpret conceptual blending as *amalgamation* (Ontañón and Plaza 2010). This has been examined for the case of $\mathcal{EL}^{++}$by Confalonieri et al. (2018). The basic idea is to generalize both input concepts gradually until a generic space, thus a (least general) generalization is reached. In the case of the houseboat-blend mentioned in the introduction, a possible generic space of "house" and "boat" would be an "object based on some foundation and with people inside". After that, the intermediate concepts are combined to determine the blend. An illustration can be seen in Figure 1 (a) and details can be found in (Confalonieri et al. 2018). Basic problems of this approach are that the generic space is not necessarily unique and thus hard to find and that it is not obvious which intermediate concepts should be combined to achieve a good blend. Therefore, we aim to enhance this amalgamation-based approach with subsymbolic information with the help of embeddings.

Knowledge Base Embeddings

Embedding information into geometric spaces is a widely used technique to employ geometric regularities for learning purposes. Especially profiting from geometric embeddings is the area of *knowledge graphs (KGs)*. Knowledge graphs are sets of $(subject, predicate, object)$-triples, e.g., $(alice, loves, bob)$. As those KGs tend to be highly incomplete, one important task is the prediction of missing triples. Here, *knowledge graph embedding (KGE)* comes into play (for an overview, see, e.g., (Ji et al. 2022)). It is based on the idea of learning an embedding (e.g., with the help of a neural network) such that instances are modeled as points and predicates (thus relations) as geometric operations in a vector space (TransE (Bordes et al. 2013) for example represents relations as vector translations). These KGE approaches can be extended to *knowledge base embeddings* by considering also background knowledge in form of an ontology (see (Bourgaux et al. 2024) for an overview). The basic idea is to add to the representation of instances and relations a representation of concepts and their interplay in a geometrical way. Therefore on the one hand the subsymbolic information of similarity of instances and concepts is used and on the other hand the symbolic information in form of an ontology. Concepts are mostly represented as some specific convex object, e.g., as boxes, e.g., by Yang, Chen, and Sattler (2025), spheres (Kulmanov et al. 2019) or convex cones (Özçep, Leemhuis, and Wolter 2020). An instance belongs to a concept if the point representing the instance lies in the the convex object representing the concept. Conjunction of concepts is mostly represented as set-intersection of the concept representations, and subsumption by the subset relation. We focus here on the DL $\mathcal{EL}_\bot$and corresponding box-embeddings, particularly on TransBox (Yang, Chen, and Sattler 2025).

Each concept A is represented as a box $box(A)$, an axis-aligned hyperrectangle in the n-dimensional space $\mathbb{R}^n$ with a center $c(A) = (c_1, \ldots, c_n) \in \mathbb{R}^n$ and an offset given by

$o(A) = (o_1, \ldots, o_n) \in \mathbb{R}^n_{\geq 0}$. Thus, $box(A)$ is defined as

$$box(A) = \{x \in \mathbb{R}^n \mid c(A) - o(A) \leq x \leq c(A) + o(A)\}$$

where $x \leq y$ when $x_i \leq y_i$ for $1 \leq i \leq n$. The conjunction of concepts is interpreted as intersection of boxes.[1] The definition of relations can be seen in Figure 1 (b). It is based on the idea of modeling relations also as boxes and states that a triple (a, r, b) holds if $a \in box(r) + b$, thus, $box(r)$ can be considered as containing reified representations of the triples. In the shown example, a box representation of the concept "Pizza" and the role "eats" is determined. The instance "alice" is represented as a point in the space. Then it is the case that "Alice eats pizza" if the point representing Alice is part of the box representing "eats pizza". This box is determined by adding up the centers resp. the offsets of the boxes of "Pizza" and "eats".

Blending Based on Box-embeddings

Assume an ontology is given and two concepts C_1 and C_2 are chosen as inputs to be blended. Based on the amalgamation-based blending approach, the two input concepts are relaxed until a common generic space is reached. This generalization is usually not unique: which part of the concepts should be relaxed and how should the generic space look like? For the houseboat-example, should, e.g., the property that a house rests on ground be relaxed, or the property that people are living inside? The embedding of the ontology offers a guiding principle for this process by offering a heuristic in form of concept similarity: when several generalizations are possible, then the one is used that 'leads towards' the other input concept.

As a first step to use this heuristic, it is necessary to create the box-embedding of a given ontology. We assume that the embedding acts as a model of the ontology, thus represents each axiom and each fact of the ontology correctly. For cases where such a perfect embedding is not found, some approximation strategies are necessary. Now, the boxes representing the two input concepts are identified.[2] The position of the two boxes relative to each other gives a first hint on the nature of the blend: if the distance between the boxes is small, then they are already quite similar and the blend will be less creative (or at least less surprising). If the distance is larger, then the concepts are quite dissimilar and the blend will be more surprising. It has been argued, e.g. by Mikolov et al. (2013) for the case of word embeddings, that similar information is embedded at similar positions. Of course, this is only possible for a rich representation of knowledge. If there are, e.g., only non-intersecting concepts and no relations, then it is not possible to infer any similarity information out of the embedding. Still, if a diverse representation is given, especially when instances are connected by relations, then, due to an enforced restricted dimensionality, it is necessary to model the information as efficiently as possible. Then, first, due to the convexity of the boxes representing superconcepts and, second, due to the limited space, it is necessary to model 'similar information' at 'similar positions' (and even predict new concepts for given instances if this eases the representation). Another interesting aspect is the negative sampling used in the embedding approaches. Negative sampling is a standard technique to improve the learning process (see, e.g.,(Qian et al. 2021) for an overview) by introducing corrupted triples for learning. Thus (subject, predicate, object)-triples inconsistent with the ontology are created which should thus not be satisfied by the embedding. This constraints the size of the boxes and prevents trivial solutions which would not contain any similarity information.

Generic Space

Given the two boxes representing the input concepts, it is straightforwardly possible to determine the generic space geometrically. By definition, a concept is a subconcept of another concept if the box representing the former concept is part of the box representing the latter concept. Therefore, a possible generic space of the two input concepts regarding the embedding can be defined as the smallest box containing both input boxes. Note that here only one interpretation is considered, and, therefore, the generalization found is not necessarily a generalization of the original ontology (due to the additionally inferred statements which are not necessary inferrable in the ontology). However, as argued by Confalonieri et al. (2018), it is not necessary to find an exact generalization but sufficient to find an approximation. In this case, the approximation is especially helpful, as it incorporates also the fact-based knowledge of the world: when, e.g., the ontology allows for many special cases not relevant for the considered use-case, then the facts do not reflect all these special cases: the embedding and therefore also the generic space can be simplified (in this case specialized) to a more specific case. Note that such a geometric representation of a generic space is always found, as a convex hull of two boxes always exists. For very dissimilar concepts, the generic space could be, however, quite large (or even trivial, thus the top-concept). Figure 2 depicts in purple the generic space GS based on the input concepts C_1 and C_2.

Blend

The generic space determines the area of possible blends. However, not every point in this space represents an instance that is a suitable blend. The problem of finding a good blend in this geometric space of possible blends can be considered as analogous to the problem of finding a good amalgamation amongst all possible ones. However, we encounter the simplification that due to the learning of the embedding, the space of possibilities has already been reduced. The amalgamation-based generalization procedure can be considered as the generation of a generalization path leading to the blend. Therefore, we interpret the search for an amalgamation as a path-finding process in the embedding space. Consider Figure 2 (a), depicting an excerpt of a two-dimensional embedding. The two input concepts and possible generalizations of the input concepts are depicted as C_i, C'_i, C''_i for $i \in \{1, 2\}$ (in real world use-cases, the

[1] TransBox offers also a different interpretation of the intersection that is not considered here.

[2] Those concepts are not necessarily represented in the embedding. This special case could be circumvented by retraining the embedding and is not further discussed here.

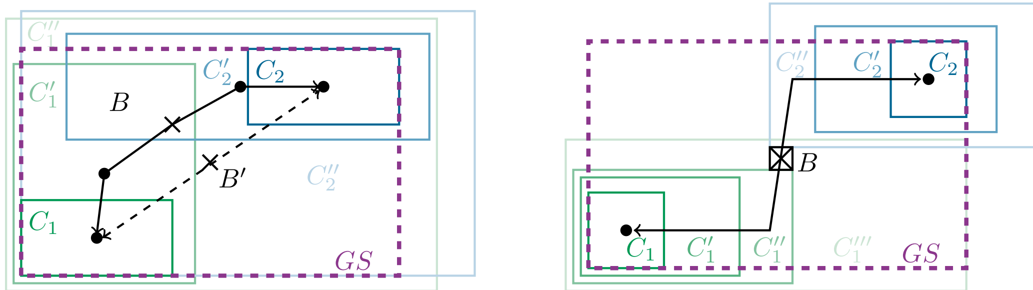

Figure 2: (a) Visualizing distinct generalization paths; (b) Visualizing two generalization paths without a direct GS-intersection

generalizations would be of course more diverse). As discussed, the blend is by definition situated somewhere in the generic space. A first attempt is to consider the Euclidean path in between the two concepts, thus a line in between the two centers of the boxes. This is depicted by blend B'. B', however, represents only a low quality blend: though it is part of the generic space, it does not model any individual features of the two input concepts and is therefore too general. This interferes with the basic principles of CB (as the blend should be a selective combination of the inputs).

Although B' could be interpreted as a blend based on amalgamation ($B' \in box(C_1'') \cap box(C_2'')$), it shows that there is some guiding principle needed to accomplish not only an arbitrary but a good amalgamation. Therefore, it is not sufficient to consider solely the shortest path. The basic idea is to interpret the amalgamation search process as a two-player game (similar to (Righetti et al. 2021)) and iteratively relax both concepts C_1 and C_2 until a common box is reached that contains B, based on a heuristic that tries to find a trade-off between specificity of the generalization and reduction of the distance to the other generalization. An example can be seen in Figure 2 (a): The two starting points of the path are the centers of the boxes representing the input concepts. Following the line of prototype theory (see, e.g., (Hampton 2006)), it is assumed that the center of a box contains instances more prototypical for a concept than instances at the border of the concept representation. Then, both input concepts are relaxed with the aim of finding a meeting point of the two paths. Thus, based on some heuristics, here C_1' and C_2' are chosen as generalizations and again the centers of the respecting boxes are chosen as new endpoints of the path. They represent prototypical representations of the more general concepts. This process can be continued until two generalizations are reached that intersect with each other. In this intersection a blend can be searched for. One possible heuristics would be to chose a blend at the boundaries of this intersection, as now not the most prototypical instance is looked for but a surprising and creative one. This depicts one variant of finding a path in the embedding space. Of course, depending on the use case, the exact heuristics and the data used, it is possible to use different strategies for determining such a path.

In this example, it was possible to determine an intersection of the generalizations after one generalization step. However, it could be necessary to consider more than one step. It could also be the case that an intersection of generalized concepts is only possible if one or both generalizations lead either to the trivial concept or enforces a too severe generalization. Then, e.g., the blend of house and boat would be a “cube-shaped object that swims”. This is due to the fact that the embedding does not model all possible combinations of concepts but only a subset of them. Considering possible blends where the generalizations are not intersecting widens the search space and especially allows for finding more surprising (as not modeled) blends. This works as follows: instead of combining generalizations to make up a blend only when the two generalizations intersect in the GS, the two generalizations are considered if they are in a sufficiently close distance. In the example of Figure 2 (b), the input concepts can be generalized to C_1'' and C_2'', respectively. These two concepts are quite similar (though not intersecting inside the GS) but still not a too broad generalization (their area has not increased too much in comparison to the input concepts). However, to gain a blend based on an intersection, it would be necessary to generalize C_1'' to C_1''', thus the resulting blend would be much broader. Therefore, here, the blend is considered as the space in between the two generalizations instead of generalizing until a common space is reached. Blends in this space are, though not directly inside the generalizations, closely related to both of them. This approach incorporates the similarity information actively, thus using that sufficient proximity in representation means we can assume that concepts C_1'' and C_2'' ‘essentially hold’ at blend B. At the same time, it is not as strict as the intersection-based approach and thus allows for more detailed and less generalized blends.

Optimality Principles

One important aspect of CB is to judge the quality of the blends based on optimality principles (Pereira and Cardoso 2003). Some of the optimality principles are already implicitly covered since the construction of the embedding provides an interpretation of the ontology and includes similarity information. Thus, e.g., the *relevance principle* should be implicitly fulfilled (stating that only relevant parts of the input concepts should be contained in the blend), as the blend is positioned in between the two inputs and therefore all concepts included in the blend are more similar to both of the input concepts than the input concepts are to each other. The *unpacking principle*, stating that it must be possible to reverse the whole blending process starting with the blend, is also respected since the construction is based on finding a path with the blend on it. Thus, to unpack we start at the blend and reverse the movement on the path towards the start- and endpoints. Another principle is about maximizing and intensifying vital relations, i.e. to prefer fundamental relations. Though such vital relations are not directly incorporated into the blend, it can be stated that the embedding is able to find important relations, as those relations also play an important role in the embedding, and are thus more likely to be also modeled in the blend. Implementing this and validating these claims experimentally (and additional criteria such as *integration* and *topology*) will be the subject of future extensions of the present work.

Experiments

In future work, we aim to develop more detailed heuristics and to explore the exact interplay of search and amalgamation. Still, to exemplify the general applicability of the approach, we here present a concise experiment based on TransBox[3] (Yang, Chen, and Sattler 2025) (introduced in the preliminaries). With TransBox, a two-dimensional embedding of a simple ontology of a horse and a bird is learned, partly displayed in Figure 3. It is based on the concept definitions used in Confalonieri et al. (2016): There, a “Horse” is defined as a “Mammal” having “Legs” and the ability to “Walk” and “Trot”. A “Bird” is an “Avialae” having “Wings” and “Legs” which is able to “Fly” and to “LayEggs”. There, a good blend is considered as being a Pegasus, thus a “Mammal” with “Legs” and “Wings” that has the ability to “Walk”, “Trot” and “Fly”. We created a simple ontology around these definitions, stating, e.g., that mammals do not lie eggs (by disregarding the platypus) and birds don’t run. Due to the simplicity of the ontology, a two-dimensional embedding is possible. In most cases, more dimensions would be needed. This simple example, though, already illustrates three vital properties of the embeddings, useful as a basis for the proposed approach. (i) In the embedding of the example ontology, it can be seen that “Animal” is actually a superclass of both “Bird” and “Horse” and, as no other more specific animal classes are considered, it constitutes a plausible generic space. The constraints are all fulfilled, e.g., horses don’t have wings and can’t fly. Moreover, it can be seen that there is a point in between the mapped input concepts, where an animal being able to walk, to trot and to fly and which has legs and wings, is represented, and thus describes Pegasus. This point can be found by generalizing the concept of “Horse” to “Mammal that can trot” and by generalizing the “Bird” to an “Animal that can fly and has wings”. Other options for generalization steps would have been, e.g., to generalize the “Bird” as an “egg-laying Avialae”, that would, however, not lead to this blend. This blending process has been carried out based on manually constructing the best blend given the embedding. As future work, a heuristic is needed to automate this process. In this simple case, the blend can be found in between the two inputs in a Euclidean sense. This is, especially when considering more complex concepts, not necessarily the case anymore (as discussed in the last section). (ii) The generic space actually depicts the space of suitable blends, thus, the blend quality decreases outside of the generic space, especially as then the resulting blend is too general. (iii) This example, finally, also shows the challenges of this approach, as even this simple ontology is modeled only approximately correctly, and thus forces us to consider techniques coping with such inaccuracies in order to approximate a good blend.

[3] https://github.com/HuiYang1997/TransBox

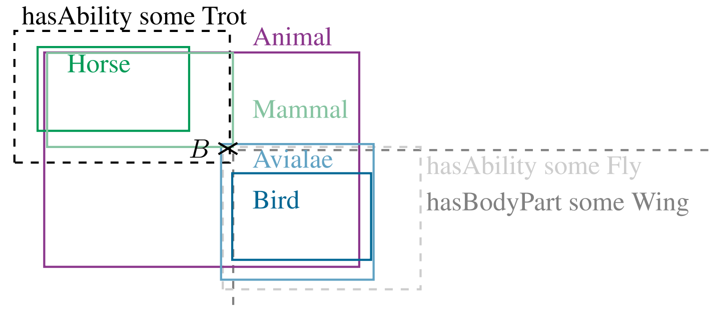

Figure 3: An example for the blend of bird and horse in the animal ontology.

Conclusion

We sketched first steps towards designing a neuro-symbolic CB framework based on the embedding of ontologies. Since existing embedding approaches suffer from a limited representational accuracy, an implemented path-finding heuristic should not only find a suitable path but should also be able to cope with partly incorrect embeddings. Additionally, considering points directly in the embedding space is a good first step to find a candidate blend. However, it restricts the space of possible blends to the ones that are already contained in the embedding space or at least approximated there. Future work is thus to enhance our approach by not only focusing on the blends in one fixed interpretation, but by either influencing the embedding process such that more and different blends can be found, or by considering the creation of the generic space only as a first step and then to use the information to find the blend in a more complex process, also including instance information. The implementation of this approach and the determination of suitable heuristics is ongoing work. In general, a more thorough examination is needed, especially for larger ontologies. However, the experiment and the discussion show that our approach promises novel neuro-symbolic strategies for performing CB based on embeddings.

Acknowledgments

We acknowledge the financial support through the ‘Abstractron’ project funded by the Autonome Provinz Bozen - Südtirol (Autonomous Province of Bolzano/Bozen) through the Research Südtirol/Alto Adige 2022 Call.

References

Baader, F.; Calvanese, D.; McGuinness, D. L.; Nardi, D.; and Patel-Schneider, P. F., eds. 2007. *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press, second edition.

Baader, F.; Brandt, S.; and Lutz, C. 2005. Pushing the $\mathcal{EL}$ Envelope. In *IJCAI'05: Proceedings of the 19th International Joint Conference on Artificial Intelligence*, 364–369.

Bordes, A.; Usunier, N.; Garcia-Duran, A.; Weston, J.; and Yakhnenko, O. 2013. Translating Embeddings for Modeling Multi-relational Data. In *NIPS'13: Proceedings of the 26th International Conference on Neural Information Processing Systems*, 2787–2795.

Bourgaux, C.; Guimarães, R.; Koudijs, R.; Lacerda, V.; and Ozaki, A. 2024. Knowledge Base Embeddings: Semantics and Theoretical Properties. In *Proceedings of the Twenty-First International Conference on Principles of Knowledge Representation and Reasoning*, KR-2024, 823–833. International Joint Conferences on Artificial Intelligence Organization.

Confalonieri, R., and Kutz, O. 2020. Blending under deconstruction: The roles of logic, ontology, and cognition in computational concept invention. *Annals of Mathematics and Artificial Intelligence* 88(5–6):479–516.

Confalonieri, R.; Schorlemmer, M.; Kutz, O.; Peñaloza, R.; Plaza, E.; and Eppe, M. 2016. Conceptual Blending in $\mathcal{EL}^{++}$. In *29th International Workshop on Description Logics, DL 2016*, volume 1577. CEUR Workshop Proceedings.

Confalonieri, R.; Eppe, M.; Schorlemmer, M.; Kutz, O.; Peñaloza, R.; and Plaza, E. 2018. Upward refinement operators for conceptual blending in the description logic $\mathcal{EL}^{++}$. *Annals of Mathematics and Artificial Intelligence* 82(1–3):69–99.

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. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.

Goguen, J. A., and Harrell, D. F. 2010. Style: A Computational and Conceptual Blending-Based Approach. In Argamon, S.; Burns, K.; and Dubnov, S., eds., *The Structure of Style*. Springer Berlin Heidelberg. 291–316.

Hampton, J. A. 2006. Concepts as prototypes. In *Psychology of Learning and Motivation Volume 46*, 79–113. Elsevier.

Ji, S.; Pan, S.; Cambria, E.; Marttinen, P.; and Yu, P. S. 2022. A Survey on Knowledge Graphs: Representation, Acquisition, and Applications. *IEEE Transactions on Neural Networks and Learning Systems* 33(2):494–514.

Kulmanov, M.; Liu-Wei, W.; Yan, Y.; and Hoehndorf, R. 2019. $\mathcal{EL}$ Embeddings: Geometric Construction of Models for the Description Logic $\mathcal{EL}^{++}$. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 6103–6109. International Joint Conferences on Artificial Intelligence Organization.

Leemhuis, M., and Kutz, O. 2024. A Cloud Full of Paths: Conceptual Blending as Betweenness Relation. In *Proceedings of The Eighth Image Schema Day (ISD8)*, volume 3888 of *CEUR Workshop Proceedings*.

Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G.; and Dean, J. 2013. Distributed Representations of Words and Phrases and Their Compositionality. *Proceedings of the 26th International Conference on Neural Information Processing Systems* 2:3111–3119.

Olearo, L.; Longari, G.; Melzi, S.; Raganato, A.; and Peñaloza, R. 2024. How to Blend Concepts in Diffusion Models. In *Proceedings of The Eighth Image Schema Day (ISD8)*, volume 3888 of *CEUR Workshop Proceedings*.

Ontañón, S., and Plaza, E. 2010. Amalgams: A Formal Approach for Combining Multiple Case Solutions. In Bichindaritz, I., and Montani, S., eds., *Case-Based Reasoning. Research and Development*. Berlin, Heidelberg: Springer Berlin Heidelberg. 257–271.

Özçep, Ö.; Leemhuis, M.; and Wolter, D. 2020. Cone Semantics for Logics with Negation. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, 1820–1826. International Joint Conferences on Artificial Intelligence Organization.

Pereira, F. C., and Cardoso, A. 2003. Optimality Principles for Conceptual Blending: A First Computational Approach. *AISB Journal* 1(4).

Pereira, F. C. 2007. *Creativity and Artificial Intelligence: A Conceptual Blending Approach*. Mouton de Gruyter.

Qian, J.; Li, G.; Atkinson, K.; and Yue, Y. 2021. Understanding Negative Sampling in Knowledge Graph Embedding. *International Journal of Artificial Intelligence & Applications* 12(1):71–81.

Righetti, G.; Porello, D.; Troquard, N.; Kutz, O.; Hedblom, M. M.; and Galliani, P. 2021. Asymmetric hybrids: Dialogues for computational concept combination. volume 344 of *Frontiers in Artificial Intelligence and Applications*, 81 – 96. IOS Press.

Wang, S.; Petridis, S.; Kwon, T.; Ma, X.; and Chilton, L. B. 2023. PopBlends: Strategies for Conceptual Blending with Large Language Models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, volume 33, 1–19. ACM.

Yang, H.; Chen, J.; and Sattler, U. 2025. TransBox: $\mathcal{EL}^{++}$-closed Ontology Embedding. In *The Web Conference 2025*.