

Does Stable Diffusion Dream of Electric Sheep?

Simone Melzi, Rafael Peñaloza and Alessandro Raganato

University of Milano-Bicocca, Milano, Italy

Abstract

Stable Diffusion is a text-to-image generation model that is based on latent diffusion. It works by first translating the textual prompt into a multidimensional *latent space*, which can be seen as an internal representation of a conceptual space. For other kinds of generative models, it has been argued that relationships between concepts can be deduced from the geometrical properties of the latent space. In this paper we explore this claim for a pre-trained Stable Diffusion model. In particular, we verify its capabilities to produce images that *blend* two concepts without any fine-tuning.

Keywords

conceptual blending, conceptual spaces, stable diffusion, generative models




Figure 1: Summary of our approach: from two separate concepts (“frog” and “lizard”) we produce the blended concept “frog lizard” through (i) interpolation in the latent space and (ii) direct prompting.

1. Introduction

Conceptual blending refers to the cognitive task of combining the properties of two different concepts to produce a new distinguished concept. A successful conceptual blend is predicated on the existence of a (potentially internal) conceptual representation where the relevant features of each concept can be identified and manipulated.

In recent years, generative methods have gained attention within the AI community. In a nutshell, generative models try to *generate* an output that is pertinent to a given input. For

The Seventh Image Schema Day, September 2, 2023, Rhodes, Greece


✉ simone.melzi@unimib.it (S. Melzi); rafael.penalozan@unimib.it (R. Peñaloza); alessandro.raganato@unimib.it (A. Raganato)

🌐 <https://sites.google.com/site/melzismn/> (S. Melzi); <https://rpenalozan.github.io/> (R. Peñaloza);

<https://raganato.github.io/> (A. Raganato)

🆔 0000-0003-2790-9591 (S. Melzi); 0000-0002-2693-5790 (R. Peñaloza); 0000-0002-7018-7515 (A. Raganato)

© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

example, text-to-image models receive a textual prompt and produce (generate) an image which visually represents the prompt. Although they differ greatly in their architecture and implementation details, most modern generative models share the same high level structure. A given input (e.g. the textual prompt) is first encoded as a highly dimensional vector, which is later decoded as the desired output medium (e.g. the image). The space where all the vectors reside—known as the *latent space*—can be thought of as a kind of *conceptual space* [1] where all the information about the concepts and their relationship is encoded. Assuming the manifold hypothesis [2], this representation should be able to encode *all* relevant concepts. In particular, each conceptual primitive has a point or region in this space. This architecture can be seen as a metaphor for cognition, where an individual e.g. reads a piece of text, producing an internal mental representation which can then be externalised in different manners like an image.

Almost since their conception, text generative models like BERT [3], GPT [4], and earlier incarnations have been analysed for their capacity to solve some cognitive tasks. A task that has gained interest for text generative models is that of analogical reasoning [5]. The question of analogical reasoning has been studied by analysing the latent space directly [6], or using the model as a black-box generator [7, 8]. In the former case, it has been argued that analogical reasoning has a correspondence with vector operations in the latent space. This means that the latent space assigns a point or a region to each primitive concept, and image schemas arise from geometric operations in this space. In other words, these results suggest that one can produce analogies simply by navigating the latent space. A natural question that arises is whether other kinds of cognitive tasks exhibit similar geometric properties; and if so, which.

Despite some clear similarities between the tasks from a cognitive point of view—both require a representation and extraction of the relevant features from a class of concepts—conceptual blending has not received the same kind of attention as analogical reasoning. This can be explained in part by the fact that conceptual blends are not easy to verify textually; that is, it is difficult to design an experiment for controlling whether a text-generative model is producing conceptual blends or not. Indeed, the externalisation of a conceptual blend is of a more graphical nature, at least for concrete concepts. An image representing the blend of two concepts will visually showcase the features of each of the original concepts.

Our goal in this paper is to verify whether conceptual blending corresponds also to vector operations on the latent space. That is, we want to see whether it is possible to produce obtain the blend of two concepts, simply by moving through the latent space, without any additional input to the model. This analysis will provide some insights into the properties of the (usually opaque) intermediate representation space, and its potential similarities to a cognitive conceptual space. To achieve our goal, we use a text-to-image generative model (specifically Stable Diffusion) to produce images that blend two different concepts through an interpolation of the two separate encodings; that is, a point in the line segment connecting them over the latent space. Our work differs from previous suggestions like [9, 10] in that we do not generate different prompts to describe the blended concept, but rather manipulate directly the latent space. As we argue, this is akin to analysing and manipulating the encoding of conceptual primitives in the internal space of the model. For comparison purposes, we also generate an image using a natural textual prompt for the conceptual blend.

As a first empirical study, we consider eight different concept pairs, selected with different criteria in mind. In particular, half of them are decompositions of common compound words,

while the other half refers to novel notions which we do not expect to observe in the training set. Although further analysis is needed, our results suggest that both, the interpolation method and a direct simple prompt, provide simple and cheap ways to obtain images of blended concepts. An analysis of the resulting outputs also provides some further insights on bias, ambiguity, and abstraction in the generative model.

Importantly, we do not blend images, but rather attempt to produce an image representing the blend of two concepts (given in a textual prompt). By our approach, we try to answer whether concept blends (and by extension, image schemas) arise naturally in models not trained to produce them. This study can shed light on the behaviour of abstract image schema representations. By the same token, we are not interested in *prompt engineering* for finding the best prompts to produce an adequate conceptual blend.

2. Stable Diffusion

We briefly introduce the main components of diffusion models and Stable Diffusion which are relevant for understanding this work. A full-fledged description is beyond the scope of this work; we refer the interested reader the main source material [11, 12].

From an abstract point of view, diffusion models belong to a class of encoder/decoder architectures which use two separate neural models. The first model (the *encoder*) translates the input into a point in the highly dimensional space \mathbb{R}^n (typically with a large n), also known as the *latent space*. The *decoder* model, on the other hand, transforms each point of the latent space into an element of the output space. Hence, for text-to-image systems like Stable Diffusion and Dall-E [13], the encoder takes as input sentences (or more in general strings) and the decoder generates an image from each latent space element.

Different approaches are distinguished by the ways they manipulate the various elements throughout the two translation steps. In particular, some modern architectures which handle text as input use a so-called *attention* mechanism, which takes advantage of the context provided by the whole sentence to disambiguate and better characterise the purpose of each word in the input during the encoding phase. What characterises diffusion models in general is that they are trained to remove the noise from a randomly generated base until the output (in our case, a picture) is obtained. Since the starting noise is randomly generated, one single (unchanged) prompt—and in particular, one single point from the latent space—may yield many different outputs. It is worth noting that this latter feature makes it difficult to systematically evaluate the behaviour of diffusion models.

One can view the latent space—that is, the intermediate, internal representation that links the input text to the output picture—of such an architecture as a *conceptual space* in which all the possible concepts recognisable by the model are encoded. Indeed, the point in the latent space that connects the input (sentence) with the output (picture) is often considered an abstract representation of the concept they refer to; see Figure 2. Intuitively, every point in the latent space represents a potentially complex concept, and nearby points are expected to represent similar notions. Thus, it is expected to satisfy the general properties of conceptual spaces *à la* Gärdenfors [1].

The idea of translating textual concepts to a latent space is not exclusive of diffusion models.




Figure 2: An example of the abstract representation of the concept of the input sentence (left) and the output picture (right) encoded by the point in the latent space (center).

It has been successfully used in natural language processing from the development of *embeddings* [14]. Already from early vector space representations of text, it was argued that the geometric properties of the latent space allow for operation-based reasoning. Mikolov et al. argue that spatial offsets can be understood as abstract relationships between concepts [6]. This insight provided the basis for performing analogical reasoning based on vector operations.

3. Conceptual Blending

Conceptual blending [15, 16] is a reasoning task in which different concepts are combined (or *blended*) to form a new concept which keeps the defining characteristics of its parts. As argued in the work first introducing it [15], blending belongs to the same class of cognitive operations as analogy, and mental modelling, among many others. It can be seen as a task of *invention* from knowledge: starting from two distinct concepts, produce one that is sufficiently distinct to be considered a new concept, but whose properties can be traced to its composing concepts.

At the moment, there is no consensus on how conceptual blending, as a cognitive task, actually works, but there is no question about the capacity of humans in performing it. It is commonly observed in comic book characters [17] or more in general in fantasy, but also as a metaphorical means to elicit certain mental images—like in the expression “sausage dog” for describing a dachshund. Linguistically, at least in English, concept blends can be produced by applying one concept as a modifier of the other. Hence a “spider man” is a man which has some characteristic borrowed from a spider. From this we can readily see that conceptual blending is not symmetric: your friendly neighbourhood *spider man* is not the same as the terrifying *man spider*. Our goal here is to verify whether concept blending capabilities can arise from the latent space of an architecture as the ones from the previous section.

It is not straightforward to analyse the availability of concept blending through text; the result tends more to be a (mental) image, an not easily verbalisable. One possible approach for showcasing blending capabilities is to use a text-to-image system. Specifically, given one such system—we focus on Stable Diffusion in this work—we want to verify whether it can produce imagery that depicts the blend of two concepts. We argue that it is easier to (subjectively) observe whether an image depicts a blended concept, than to do so for a long textual description. Importantly, we consider Stable Diffusion *as is*, with no fine-tuning or further training, to




Figure 3: The two blending approaches for the concept “snake horse:” the blended prompt in blue and the interpolation approach from red to green.

analyse its intrinsic capabilities.

For comparison purposes, we consider two approaches to concept blending in Stable Diffusion. The first approach takes advantage of the linguistic capabilities of the text-to-latent space encoder, and provides the prompt (henceforth called the *blended prompt*) which describes the blend. Hence, for instance, to produce a blend of a snake and a horse, we introduce the blended prompt “snake horse.” If the latent space of Stable Diffusion has geometric properties akin to those observed in language models, then it should be possible to blend two concepts through shifts in the latent space itself. Intuitively, each concept should be represented by a *region* in the latent space. When the encoder maps one prompt to a point p_1 in the latent space, all points that are close to p_1 should represent similar notions, which become more distinct as the distance increases. Thus, as we move from p_1 to the encoding p_2 of a second prompt, the concept shifts from the first to the second prompt. This motivates the second approach.

Our second method is based on interpolation in the latent space: if p_1 and p_2 are the latent space encodings of concepts C_1 and C_2 , then all the points in the line connecting p_1 and p_2 represent blends of C_1 and C_2 giving more or less weight to each of the original concepts. For the scope of this paper we consider the point exactly midway between the two encodings and, to go beyond the symmetrical blending, also the points at $\frac{1}{4}$ and $\frac{3}{4}$ of the way. Intuitively, if the prompts are “snake” and “horse,” the three interpolation points should construct a “snake horse” (25% snake, 75% horse), a “horse snake,” and a mix half horse and half-snake. A comparison of both approaches is depicted in Figure 3. Note that by the nature of the encoder, the latent space point corresponding to the blended prompt may be quite far away from the two individual prompts being blended. We emphasise once again that we are *not* interested in blending two images, but rather in producing a picture representation of the blend of two concepts. The limit images (in the example “snake” and “horse”) are only provided as reference points. Our input is not these images, but rather the textual prompts, which we use to analyse the capabilities of the latent space to produce conceptual blends.

Table 1

Eight pairs of concepts for testing blending capabilities

Compound Words			Novel Blends		
Concept 1	Concept 2	Blended Prompt	Concept 1	Concept 2	Blended Prompt
butter	fly	butter fly	bumblebee	lion	bumblebee lion
jelly	fish	jelly fish	elephant	duck	elephant duck
rain	bow	rain bow	frog	lizard	frog lizard
spider	man	spider man	snake	horse	snake horse

4. Imagine an Elephant Duck

In order to test the capacity of Stable Diffusion to produce blends through the two approaches described before, we constructed the blends for eight different pairs of concepts. The choice of the concepts used responds to several needs. First of all, to avoid artefacts caused by complex prompting, each concept should be describable with a single word. Second, all concepts must be concrete, with intuitive graphical representations. Third, the blended prompt must conjure a plausible (if not necessarily existing) conceptual image.

In addition to these constraints, we wanted to verify that Stable Diffusion had not previously “learned” the blended concept, and see how it behaves when it has. Thus, four pairs of words were constructed separating common compound words into their components, and the remaining four were chosen from previous blending attempts [18] and popular culture aiming for seldom (if at all) represented imagery. The eight pairs are shown in Table 1. Note that the blended prompts for compound words use their elements as distinct entities, and hence we ask for e.g. a “jelly fish” rather than a “jellyfish.” An adequate blend should produce a *fish with the properties (or made) of jelly*, rather than the well known medusozoön.

We generate all images using Stable Diffusion 2-1 [11].¹ For each of the eight pairs, we generate six images: one corresponding to the blended prompt, and five corresponding to the interpolation from the first individual prompt (e.g. “snake”) to the second prompt (“horse” in the example) at 25% intervals as in Figure 3. Recall that in diffusion models, the decoding phase going from the latent space to the generated image depends on a random noise initialisation. To avoid differences caused by the random noise generator, we fix a common seed for all six images. In this way, changes are attributable to diffusion (i.e., the decoding step) and not to the random initialization. The risk of fixing a seed is that the quality of the resulting blend may depend on the seed choice. We deal with this issue generating 10 different sets of images for each pair, each one with a different random seed. Thus, overall we produce 60 images for each concept pair, for a total of 480 pictures in the experiment. All figures are publicly accessible at <https://git-ricerca.unimib.it/rafael.penalozanyssen/isd7-images/>.

4.1. Novel Blends

Figure 4 shows the results for the blended prompt “elephant duck.” As it can be seen, the quality of the resulting blends is highly variable. Considering the exact mid-point in the interpolation (third column) it is possible to observe cases which can be thought of as “blends.” In particular,

¹<https://huggingface.co/stabilityai/stable-diffusion-2-1>




Figure 4: Two different blending solutions on the example pair (duck, elephant). On the left, the interpolation in three steps from duck to elephant representation in the latent space. The duck weight decreases as (1, 0.75, 0.5, 0.25, 0) while the elephant weight grows in the opposite way. On the right, the examples generated blending the textual prompt as "elephant duck".



Figure 5: Interpolation from “butter” (left) to “fly” (right), through a middle point depicting a butterfly.

the rows 3, 7, and 10 showcase animals combining the properties of elephants and ducks. The other partial interpolations (second and fourth columns) on the other hand mainly represent the limit concept (a duck, or an elephant, respectively) except for rows 7 and 10, where one could plausibly interpret the image as an elephant duck. On the other hand, the blended prompt (last column) generates several positive instances, but has the issue that “elephant” dominates over “duck” (despite duck being the main noun in the prompt). Nonetheless, the results do showcase blended concepts.

Not all pairs show this kind of behaviour. The blended prompt for “bumblebee lion” produces a clearly distinguishable lion; typically a lion face. In general, also the middle interpolation produces lions, although in this case a few more bumblebee features become visible—albeit, when explicitly searching for them. Regardless, the concept lion seems to be dominant w.r.t. bumblebee. This gives us some insights about the relative weight of primitive concepts within the latent space.

4.2. Compound Words

Unsurprisingly, the blended prompt for pairs derived from compound words produce images of the compound word (despite the space separating the components) rather than the expected blend. This behaviour showcases the presence of bias in the training data, where the blended prompt word is likely considered a misspelling of the original compound word. An extreme example is given by the pair “spider” “man,” where the blended prompt invariably produces the well-known comic-book superhero dressed in red and blue. Extreme bias is apparent from knowledge that neither spiders nor men are associated to the bright coloring of Spider-man, which only exists due to printing limitations and artistic whims.

According to the geometric interpretation of the latent space, the interpolation-based approach should bypass this bias, and produce a higher variability in the blends. In general, the interpolation approach did not produce any clearly recognisable blends, with the only exception being some conceptual fish imagery which may be interpreted as being made of jelly. We speculate that this could be partially caused by the choice of relatively abstract or ambiguous concepts in the pairs. This can be seen by analysing the pictures generated for the concept “bow” which include a knot, the action of bowing, but also nature scenery.

We found a surprising result through interpolation, though. Figure 5 depicts one full interpolation chain from “butter” (left) to “fly” (right). The interpolation makes a natural transition from a depiction of butter, to a depiction of a fly while, interestingly, producing a depiction of a

butterfly in the process. We believe that it is worth investigating this behaviour further.

4.3. Further Insights

A compulsory analysis of the generated images showcases bias in unexpected places. We have mentioned the not-so-amazing spider man results. More surprising are the results regarding the prompt “bumblebee.” Where one may expect a chubby stripped insect, Stable Diffusion returns, in 10 out of 10 calls, the modular self-configuring autobot from the *Transformers* franchise. Figure 6 shows six of those ten results; the remaining four fall into the same class.

From a different perspective, we find that Stable Diffusion struggles with concepts which are too general or abstract. One example was the notion of bow. Another one is the concept “man” for which the model produces everything going from a crowd, to a street, to a city, to a grid of houses (along two pictures classifiable as men). In hindsight, we should have expected these results, as “man” can have many different interpretations depending on the context.

5. Conclusions

We studied the capacity of Stable Diffusion to generate blended concepts through two techniques: one by making a direct prompt, and the other by manipulating the points in the latent space. Our work uses a pre-trained Stable Diffusion model and requires no additional training or fine-tuning. The results are promising, although they still leave some space for improvement. In particular, the output evaluation relies on a subjective observation of the generated images. From a conceptual space point of view, our results provide evidence that the latent space is capable of encoding primitive concepts, and manipulations in this space provide the basis for more complex image schemas.

The blended prompt approach is limited by the impact of bias from the training corpus. This is particularly relevant given the simplicity of the prompts that we chose, where each concept is described by a single word. More complex and detailed prompts could alleviate this issue. This idea was explored in [9], where a large language model (LLM) is used to generate prompts. We chose not to follow this path because first, it becomes difficult to separate the capacity of the LLM to write a good prompt from the capacity of Stable Diffusion to generate a good blend; and second, we are more interested in the properties of the latent space, and the blended prompt was chosen for comparison purposes.

For the interpolation method, we used three intermediate points to generate the images. The results suggest that moving only one quarter of the way from a latent space point to another



Figure 6: Six images for the prompt “bumblebee.”

does not produce noticeable changes in general. It may be interesting to analyse how the behaviour changes at different interpolation distances. From our observations, the latent space seems to partially encode conceptual blends without any special training. It would be interesting to analyse the possibilities of using blends during training to improve the latent representation. Another avenue for future research is to analyse different geometric or space manipulation methods for conceptual blending beyond the two presented here. To emphasise, the Stable Diffusion generator model has the scope of understanding graphically the abstract concept that is encoded at different points of the latent space.

Our view on conceptual blending does not follow the classical view by Fauconnier and Turner [15], where different spaces are constructed. The reason is that our goals differ. While Fauconnier and Turner try to explain *how* are conceptual blends constructed (and how they may be generated automatically), we are only exploring *whether* blended concepts are represented in the latent space at a position discoverable through simple geometric manipulations.

Acknowledgments

This work was partially supported by the MUR for the Department of Excellence DISCo at the University of Milano-Bicocca and under the PRIN project PINPOINT Prot. 2020FNEB27, CUP H45E21000210001; and by the NVIDIA Corporation with the RTX A5000 GPUs granted through the Academic Hardware Grant Program to the University of Milano-Bicocca for the project “Learned representations for implicit binary operations on real-world 2D-3D data.” The authors also wish to acknowledge CSC-IT Center for Science, Finland, for computational resources provided.

References

- [1] P. Gärdenfors, *Conceptual spaces - the geometry of thought*, MIT Press, 2000.
- [2] B. C. A. Brown, A. L. Caterini, B. L. Ross, J. C. Cresswell, G. Loaiza-Ganem, The union of manifolds hypothesis and its implications for deep generative modelling, *CoRR abs/2207.02862* (2022). doi:[10.48550/arXiv.2207.02862](https://doi.org/10.48550/arXiv.2207.02862).
- [3] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional transformers for language understanding, *CoRR abs/1810.04805* (2018). URL: <http://arxiv.org/abs/1810.04805>.
- [4] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, Improving language understanding with unsupervised learning (2018). URL: <https://openai.com/research/language-unsupervised>.
- [5] M. Mitchell, Abstraction and analogy-making in artificial intelligence, *Ann. of the New York Ac. of Sciences* 1505 (2021) 79–101. doi:<https://doi.org/10.1111/nyas.14619>.
- [6] T. Mikolov, W.-t. Yih, G. Zweig, Linguistic regularities in continuous space word representations, in: *Proc. of the 2013 Conf. of the North American ACL: Human Language Technologies*, ACL, 2013, pp. 746–751. URL: <https://aclanthology.org/N13-1090>.
- [7] T. Webb, K. J. Holyoak, H. Lu, Emergent analogical reasoning in large language models, 2023. *arXiv:2212.09196*.

- [8] A. Ushio, L. Espinosa Anke, S. Schockaert, J. Camacho-Collados, BERT is to NLP what AlexNet is to CV: Can pre-trained language models identify analogies?, in: Proc. of the 59th Annual Meeting of the ACL (Volume 1: Long Papers), ACL, 2021, pp. 3609–3624. doi:10.18653/v1/2021.acl-long.280.
- [9] S. Ge, D. Parikh, Visual conceptual blending with large-scale language and vision models, in: Proc. of the 12th Intern. Conf. on Computational Creativity, ACC, 2021, pp. 6–10. URL: https://computationalcreativity.net/iccc21/wp-content/uploads/2021/09/ICCC_2021_paper_90.pdf.
- [10] S. Wang, S. Petridis, T. Kwon, X. Ma, L. B. Chilton, Popblends: Strategies for conceptual blending with large language models, in: Proc. of 2023 CHI Conference on Human Factors in Computing Systems, CHI '23, ACM, 2023. doi:10.1145/3544548.3580948.
- [11] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, in: IEEE/CVF Conf. on Computer Vision and Pattern Recognition, CVPR 2022, IEEE, 2022, pp. 10674–10685. doi:10.1109/CVPR52688.2022.01042.
- [12] J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, S. Ganguli, Deep unsupervised learning using nonequilibrium thermodynamics, in: Proc. of the 32nd Intern. Conf. on ML, ICML 2015, volume 37 of *JMLR Workshop and Conf. Proceedings*, JMLR.org, 2015, pp. 2256–2265. URL: <http://proceedings.mlr.press/v37/sohl-dickstein15.html>.
- [13] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, I. Sutskever, Zero-shot text-to-image generation, in: Proc. of the 38th Intern. Conf. on Machine Learning, ICML 2021, volume 139 of *Proceedings of Machine Learning Research*, PMLR, 2021, pp. 8821–8831. URL: <http://proceedings.mlr.press/v139/ramesh21a.html>.
- [14] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, in: Proc. of the 27th Annual Conf. on Neural Information Processing Systems, 2013, pp. 3111–3119. URL: <https://proceedings.neurips.cc/paper/2013/hash/9aa42b31882ec039965f3c4923ce901b-Abstract.html>.
- [15] G. Fauconnier, M. Turner, Conceptual integration networks, *Cognitive Science* 22 (1998) 133–187. doi:10.1207/s15516709cog2202_1.
- [16] L. D. Ritchie, Lost in “conceptual space”: Metaphors of conceptual integration, *Metaphor and Symbol* 19 (2004) 31–50. doi:10.1207/S15327868MS1901_2.
- [17] G. Guizzardi, R. Peñaloza, M. M. Hedblom, O. Kutz, Under the super-suit: What superheroes can reveal about inherited properties in conceptual blending, in: Proc. of the 9th Intern. Conf. on Computational Creativity, ACC, 2018, pp. 216–223. URL: http://computationalcreativity.net/iccc2018/sites/default/files/papers/ICCC_2018_paper_56.pdf.
- [18] P. Martins, T. Urbancic, S. Pollak, N. Lavrac, A. Cardoso, The good, the bad, and the AHA! blends, in: Proc. of the 6th Intern. Conf. on Computational Creativity, ICC 2015, computationalcreativity.net, 2015, pp. 166–173. URL: http://computationalcreativity.net/iccc2015/proceedings/7_3Martins.pdf.