

Typographers in the Loop*

Mapping the Typographic Latent Space of Digits as a Matter of Responsible Design

Alayt Issak^{1,†}, Sarthak Kakkar¹, Sair Goetz², Nik Brown¹ and Casper Hartevelde¹

²Independent Typographer

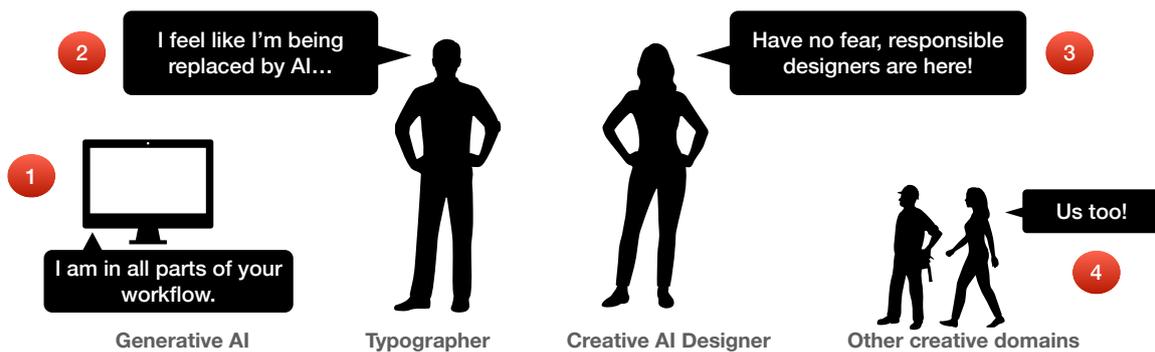
¹Northeastern University, Boston, MA 02120, USA

Abstract

Since the advancement of handwritten text to typefaces on a computer, the human mind has evolved towards corresponding various typefaces as norms of comprehension. Current-day typefaces, much like those written by hand, exist in disparities and are governed by consensus reached among Typographers. Currently, the PANOSE system, developed in 1998, is the most widely used and accepted method for classifying typefaces based on 10 visual attributes. In this work, we employ Disentangled Beta-VAE's, in an unsupervised learning approach, to map the latent feature space with a dataset of MNIST Style Typographic Images (TMNIST-Digit) of 0-9 digits across 2990 unique font styles. We expose the learning representation across a variety of font styles to enable typographers to contemplate and identify new attributes to their classification system. We also exercise AI in such a manner to promote responsible design practices.

Keywords

Typography, Latent space, Responsible design, Computational Design



History

Typography has been evident since hieroglyphics and replicating text via moveable type has also existed via woodblock printing during the Tang Dynasty (618–907 AD) in China (Figure 1) [2]. In the Latin Type, typefaces emerged as calligraphy was overtaken by the incentive for mass production, and in the history of Western printing, was inaugurated by the letterpress in 1450 [3].

Yet in its ontology, Latin typefaces are mostly named after the people that devised them, such as Garamond in France, or the movements that inspired them, such as Roman (the foreground to Times New Roman) in Italy [4].

Computational Design and Computer-aided Creativity ICCV'25 Workshop (Campinas, BR & Online)

*This pictorial is a revised, visual and extended version (implications for responsible design) of an already published paper at the ICLR 2023 Tiny papers [1].

†Corresponding author

✉ issak.a@northeastern.edu (A. Issak); kakkar.sa@northeastern.edu (S. Kakkar); sairgoetz@gmail.com (S. Goetz); nik.brown@northeastern.edu (N. Brown); c.hartevelde@northeastern.edu (C. Hartevelde)

🆔 0009-0005-8149-7358 (A. Issak); 0000-0002-8164-0303 (C. Hartevelde)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

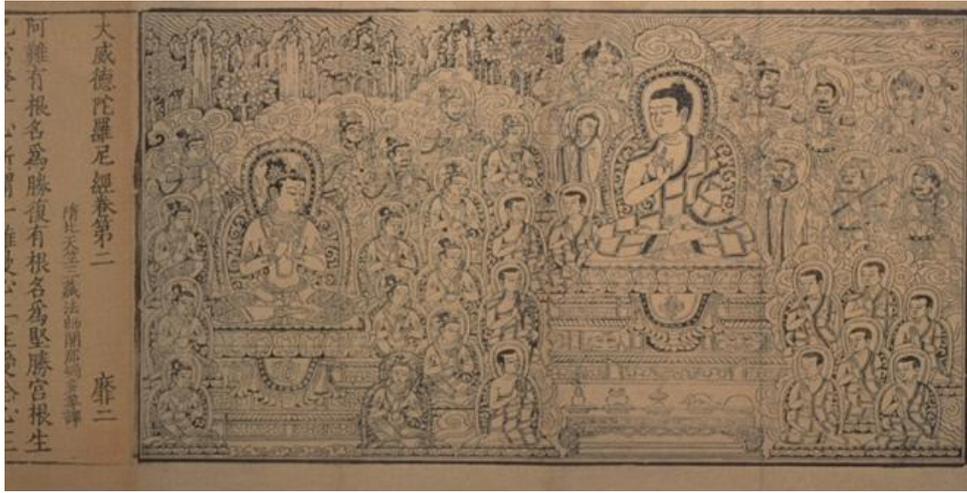


Figure 1: Chinese woodblock printing

Proceeding into later formalization, a serif (a projection on the stroke of a letter) was then introduced to mitigate the confusion one might have between differentiating nuances, such as the number ‘1’ from the letter ‘l’, and existing mechanisms that classify typefaces into five categories present as follows – Serif, Sans-Serif (without serif), Script, Decorative and Blackletter [5] (Figure 2).

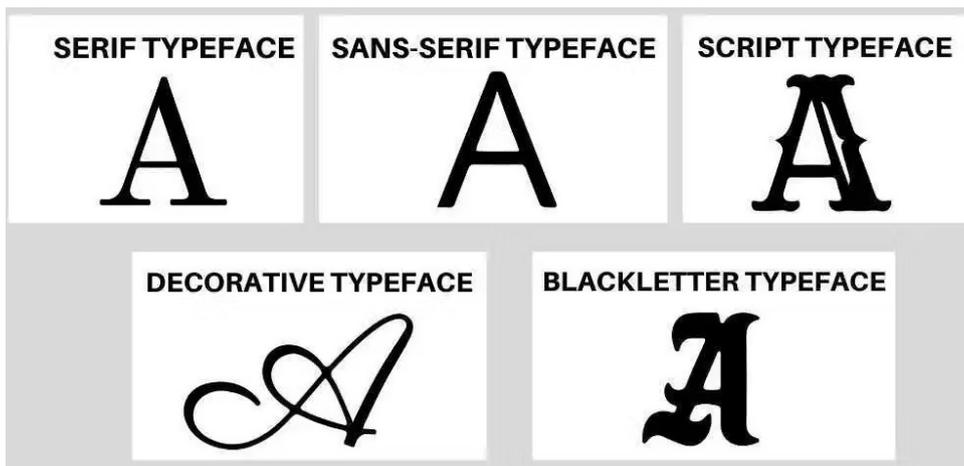


Figure 2: Typeface Classifications

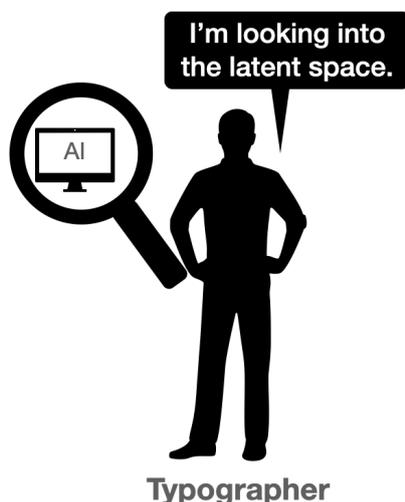
Typeface Gap

Of importance, while engaging with the typographer in our team, we found that the exact consensus to the patterns observed from the series of heuristics and matriculation is unknown and lacks comprehension. This is especially important as PANOSE, the current industry standard for classifying typefaces based on 10 attributes [6], inherits the aforementioned historical design process with downstream discrepancies.

Garnering an example on the legibility of fonts, one study found Helvetica, a **sans-serif** font type, to be successful for readers with dyslexia [7], whereas another study found **serifs** to be quite encouraging [8]. Elsewhere, notable typographers such as Sofie Beier have signaled that typeface legibility varies significantly depending on the reading situation [9], leading the overarching nature of discovery into question.

Recentering the Creative Practitioner

Recentering the creative practitioner into the nature of discovery, we seek to expose the typographic latent space of an encoder to represent typeface letter forms so typographers may identify attributes separate from those outlined by PANOSE. As such, we identify the nascent needs of typographers with the integration of AI into their creative domain. We also present AI in such a manner to recenter the creative practitioner in utilizing the technology, alongside Human-Centered AI initiatives [10].



Articulating our methods, we utilize Disentangled Beta-Variational AutoEncoders (β -VAE) as our architecture to generate typefaces across a range of learned attributes due to its robust learning of disentangled representations [11]. For our dataset, we utilize the TMNIST-Digits dataset [12]—a custom dataset consisting of 29,900 examples with 10 digits of MNIST-style images (0-9) for each of the 2900 font styles.

Programming our model, we split our dataset into a 80, 10, 10 train, validation, and test split, and train our model with the Pythae Disentangled Beta-VAE library unified by Chadebec et al. [13] for 20 epochs with a variation of hyperparameters. We utilize Adam to update our weights and employ early stopping with a patience level of 5. Upon obtaining 16 as our optimal dimension for the latent space (similar to MNIST), below, we visualize samples from a Gaussian Mixed Model and Standard Normal distribution of the latent space learned by the Disentangled β -VAE. We also present a reconstruction of a ground truth image to check our fidelity. The source code is publicly available on GitHub¹.

VAE Latent Space

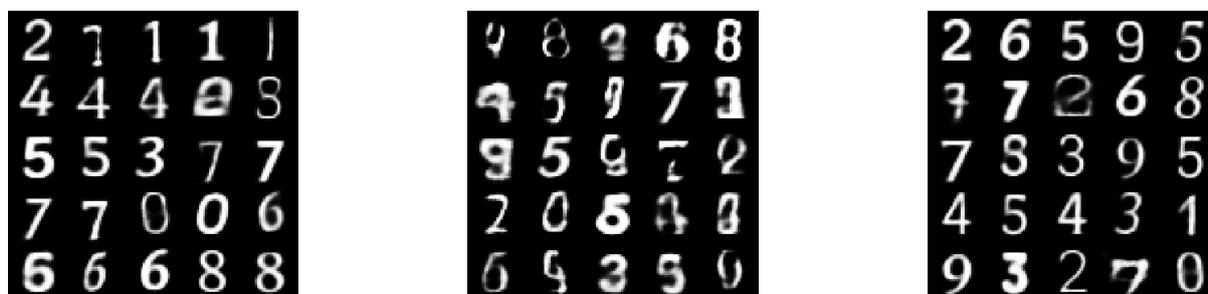


Figure 3: Sampling the Latent Space via Gaussian Mixed Model (left) and Normal Distribution (center), along with Reconstruction on the test set (right) with hyperparameters: $\beta = 2.5$, $C = 25$, $epoch = 15$.

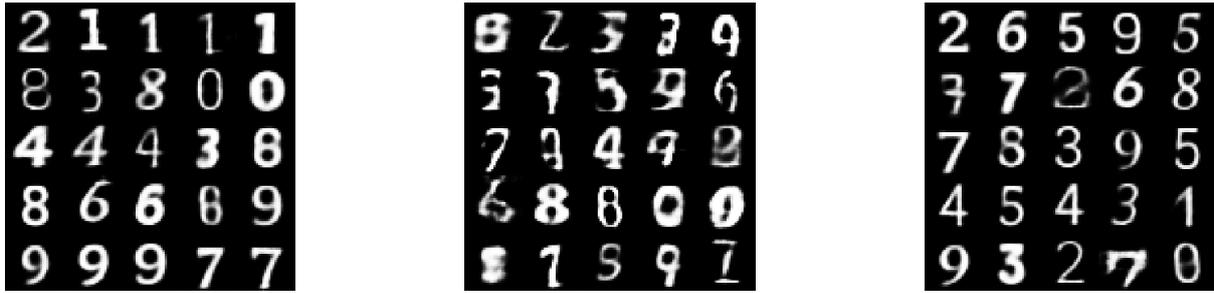


Figure 4: Sampling the Latent Space via Gaussian Mixed Model (left) and Normal Distribution (center), along with Reconstruction on the test set (right) with hyperparameters: $\beta = 4.0$, $C = 50$, $epoch = 30$.

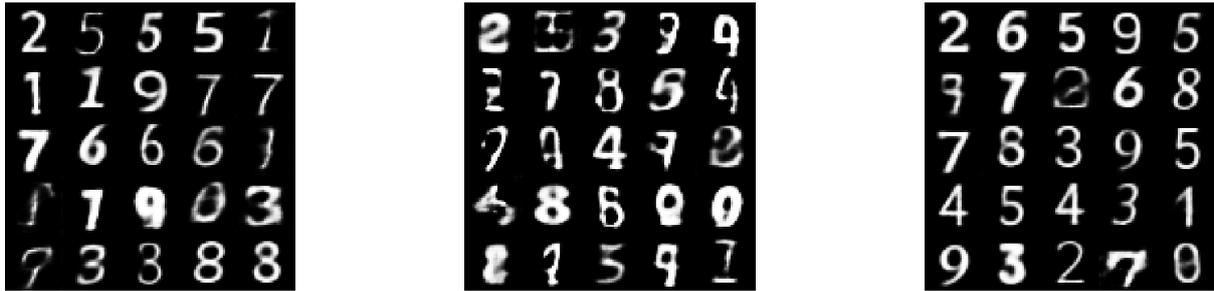


Figure 5: Sampling the Latent Space via Gaussian Mixed Model (left) and Normal Distribution (center), along with Reconstruction on the test set (right) with hyperparameters: $\beta = 5.0$, $C = 75$, $epoch = 50$.

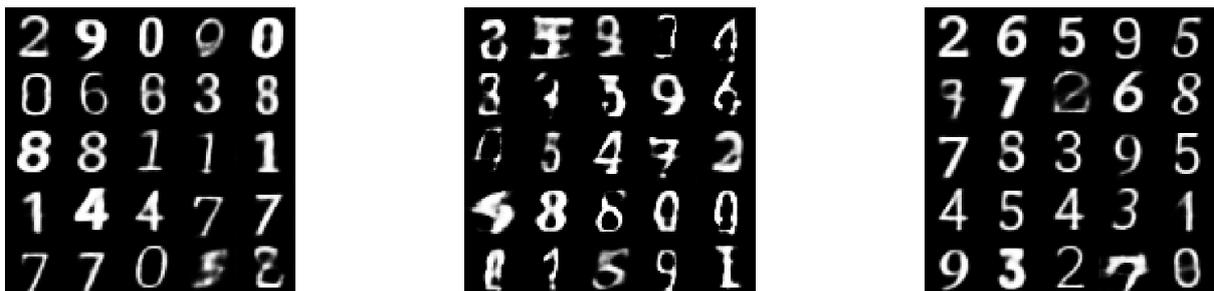


Figure 6: Sampling the Latent Space via Gaussian Mixed Model (left) and Normal Distribution (center), along with Reconstruction on the test set (right) with hyperparameters: $\beta = 6.0$, $C = 100$, $epoch = 50$.

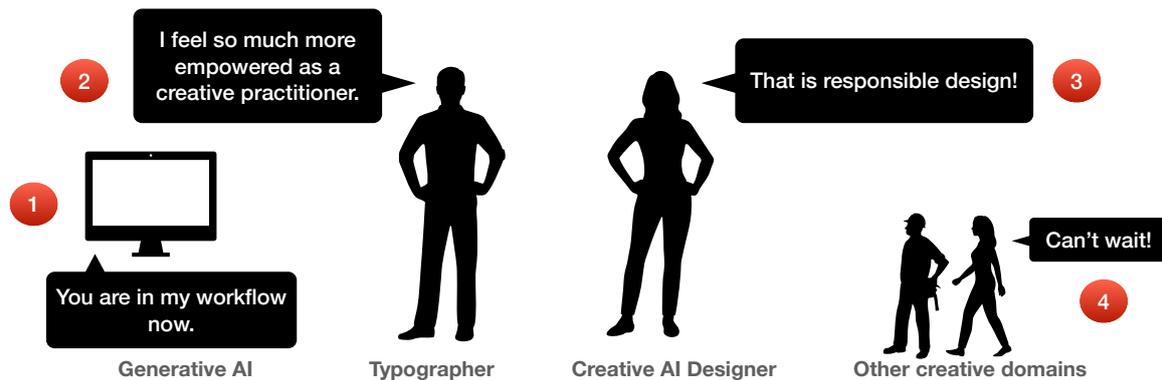
Highlighted in Figures 3, 4, 5 and 6 above, the latent space exhibits varying fidelity in our model's ability to encode representations across the 10 digits. In Figure 3 (Gaussian), the digits 8 and 4 seem to learn well whereas 0, 2, and 9 seem to be intertwined as per the poor reconstruction in the rightmost column of Figure 3. They also lose encoding in the latent space of both samples across most figures. For example, view the loss in Figures 5 and 6. Of the sampling methods, Normal Distribution also seems to sample better than a Gaussian Mixed Method, leading it to be a likely method of invoking the latent space. Thus, as per the deduction of letter forms extracted within these samples, we present this analysis for typographers to investigate current heuristics, identify novel font types, and accordingly, attribute distinctions within the latent space of fonts.

Overall

Yielded in this study, we find that mapping the "learning" that goes on in a model is a method to reverse the current heuristics by presenting the features the model has learned. Likewise, we present this endeavor to typographers so they may use these findings as a representation for unmasking

¹Github repo: <https://github.com/Sarthak-Kakkar-03/Typographic-Latent-Space>

current attributes of industry-standard font classification and matching systems, and to empower their creative practice given the burgeoning integration of AI within the creative domain.



Going Forward

We seek to extend our findings toward letters (alphabets), non-Latin typefaces, and the overarching question of typeface legibility.

Acknowledgements

We thank Nimish Magre for the insightful comments, guidance, and curating the TMNIST dataset.

References

- [1] A. Issak, S. Kakkar, S. Goetz, N. Brown, C. Harteveld, Mapping the Typographic Latent Space of Digits, 2023. URL: <https://openreview.net/forum?id=ufA2FuCGyz>.
- [2] A. A. Museum, The Invention of Woodblock Printing in the Tang (618–906) and Song (960–1279) Dynasties | Education | Asian Art Museum, 2024. URL: <https://education.asianart.org/resources/the-invention-of-woodblock-printing-in-the-tang-and-song-dynasties/>.
- [3] F. Forrest, The history of typography: retracing letter forms from hieroglyphs to Helvetica, 2021. URL: <https://99designs.com/blog/design-history-movements/the-history-of-typography/>, section: Design history & movements.
- [4] J. M. Wells, Typography - History of typography | Britannica, 2022. URL: <https://www.britannica.com/technology/typography/History-of-typography>.
- [5] B. Priya, The Classification of Typeface Styles: The Definitive Guide, 2019. URL: <https://www.turing.com/kb/classification-of-typeface-styles>.
- [6] R. Stevahn, PANOSE on the Web, 1996. URL: <https://www.w3.org/Printing/stevahn.html>.
- [7] L. Rello, R. Baeza-Yates, Good fonts for dyslexia, in: Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility, ACM, Bellevue Washington, 2013, pp. 1–8. URL: <https://dl.acm.org/doi/10.1145/2513383.2513447>. doi:10.1145/2513383.2513447.
- [8] R. Hillier, Sylexiad. a typeface for the adult dyslexic reader, Journal of Writing in Creative Practice 1 (2008) 275–291. URL: https://intellectdiscover.com/content/journals/10.1386/jwcp.1.3.275_1. doi:https://doi.org/10.1386/jwcp.1.3.275_1.
- [9] S. Beier, J.-B. Bernard, e. Castet, Numeral legibility and visual complexity, 2018. doi:10.21606/drs.2018.246.
- [10] B. Shneiderman, Human-Centered AI, Oxford University Press, Oxford, 2022.
- [11] C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, A. Lerchner, Understanding disentangling in β -vae, 2018. URL: <https://arxiv.org/abs/1804.03599>.

- [12] N. Magre, N. Brown, Typography-MNIST (TMNIST): an MNIST-Style Image Dataset to Categorize Glyphs and Font-Styles, 2022. URL: <http://arxiv.org/abs/2202.08112>, arXiv:2202.08112 [cs].
- [13] C. Chadebec, L. J. Vincent, S. Allasonnière, Pythae: Unifying generative autoencoders in python – a benchmarking use case, arXiv preprint arXiv:2206.08309 (2022). URL: <https://arxiv.org/abs/2206.08309>.