xCoAx 2021 9th Conference on
Computation, Communication, Aesthetics & X

2021.xCoAx.org

Creating Stylised Geographic Maps with Neural Style Transfer

Keywords: Computational Design, Map Art, Map Stylisation, Neural Style Transfer

André Santos
alsantos@student.dei.uc.pt
University of Coimbra, Department of Informatics Engineering, Portugal

Tiago Martins
tiagofm@dei.uc.pt
University of Coimbra, Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Portugal

João Correia
jncor@dei.uc.pt
University of Coimbra, Centre for Informatics and Systems of the University of Coimbra, Department of Informatics Engineering, Portugal

Despite regular navigational applications, geographic maps possess a strong aesthetic dimension that can be entirely separated from utilitarian tasks and purposes. We present a computational system that uses neural style transfer and open-access geographic data to generate stylised maps from any region in the world customised in the style of any given image. We showcase and analyse output maps generated by the system while offering insight into how changes to the inputs can result in clearer and more unique outputs. Finally, the real-world applicability of the output maps is addressed, leading us to conclude that they are a viable way to aesthetically establish connections to geographic places, which can happen through their application in purely decorative or more meaningful design contexts.

1. Introduction

The geographic map holds a strong historical connection to the visual arts and, throughout its extensive existence, has fulfilled both utilitarian and aesthetic needs (Rees 1980). Before science and technology reigned over mapmaking, the lack of geographic knowledge and cartographic expertise placed map creation at the hands of those proficient in the arts (i.e. people with the necessary drawing and painting skill required to translate uncertain terrain into visual elements), which is one of the reasons why, at the time, the map gained traction as a form of decorative art (Rees 1980). Centuries later, the post-digital revolution rise and development of the computer replaced the need for manual labour with autonomous digital rendering based on precise geographic data (Thrower 2008). Artistic proficiency was no longer as essential to the process of map creation as it once was, but that did not mean aesthetic interest ceased to exist. In fact, by increasing the accessibility and triviality of map creation, the computer brought about a new digitally supported interest in both the artistic and aesthetic dimensions of the map (Caquard, Piatti, and Cartwright 200; Kent 2017). These dimensions can be entirely separated from the more scientific and functional aspects, this time not due to a lack of knowledge or expertise but as a deliberate way of enhancing and focusing on the aesthetic experience of looking at a map (Kent 2017).

Aesthetically driven maps warrant visual exploration and customisation as those are the processes that lead to new creative methods of translating space into a visual medium. Artificial Intelligence (AI) presents techniques capable of assisting artists and designers in developing and experimenting with new visual styles, notably neural style transfer allows the artistic style from one image to be transferred onto another (So 2018). When used for map stylisation, this technique provides a practical way of generating a nearly limitless number of highly diverse stylised maps where the otherwise complex parameters that control the look of the output are compressed into a simple and intuitive visual input — the style input, an image file.

In this paper, we present a system[1] focused on generating stylised geographic maps through the use of a neural style transfer technique. We use open-access geographic data from OpenStreetMap[2] and other open data projects to render input maps, and an existing implementation of arbitrary neural style transfer (arbitrary meaning we can use any style image to generate outputs) to style them based on a separate input image. The output maps are generated in image format (PNG) and can serve a myriad of purposes, from purely decorative to

1. The system presented in this paper is documented online at https://cdv.dei.uc.pt/stylised-maps/

2. https://www.openstreetmap.org/#map=7/39.602/-7.839

more meaningful design objects. The driving force behind our system is creating map-like artefacts with visual interest that can fulfil aesthetic needs, thus, as it pertains to this paper, spatial accuracy and traditional communication-focused mapmaking rules are secondary to visual quality and appeal.

The remainder of this paper is organised as follows. Section 2 describes work relating to map stylisation using neural style transfer and similar artificial intelligence techniques. Section 3 provides a brief overview of our approach to develop the presented system. Section 4 showcases results and analyses their aesthetic quality and real-world applicability. Finally, Section 5 presents final conclusions and directions for future work.

2. Related Work

The primary domain of the presented system is the use of neural style transfer as a method of map stylisation. First proposed by Gatys, Ecker, and Bethge (2015), neural style transfer is a technique that uses convolutional neural networks to generate an output image that combines the semantic content of one image with the style of another. With it being a relatively recent technique, neural style transfer has yet to be fully realised in the context of geographic map stylisation. Nevertheless, while not entirely related, some projects do share some aspects, either in technique or style of output, to our own and will be presented in the following paragraphs.

Bogucka and Meng (2019) used neural style transfer to assess its capability in transferring emotions from individual paintings to a map. A group of paintings created by Berlin residents tasked with visually expressing their emotions about the city were used to create a group of stylised maps of the same place. The authors ran a survey and concluded that some emotions were successfully transferred from painting to stylised map. The attention-grabbing potential of the outputs was also noted by the authors. Although this aspect has not been explored in the context of geographic maps, there are some examples of its use to style aerial photographs. Some members of Consilium Technology (2018) developed a one-day project motivated by the team wondering how the great painters of the past would have represented the earth from an aerial top-down view. This question led them to experiment with an implementation of neural style transfer that was used to stylistically represent satellite views. Similarly, but with more selective criteria for choosing the input images, Morris (2016) managed to produce images with added visual interest using the same technique. Morris combined satellite images, mainly from NASA, that were already visually captivating by themselves with abstract expressionist paintings, resulting in colourful and unique renditions of interesting geographic locations.

While technically not neural style transfer, other projects have used similar artificial intelligence techniques to produce stylised images based on maps. OpenDot Lab trained individual neural network models to be able to recreate a satellite view of a city from just its map. After having three models trained on different cities (Los Angeles, Milan and Venice), the team was able to feed other maps to the models and thereby generate stylised images that combine the urban layout of the input map with the top-down visual characteristics of the model's city (Kogan et al. 2016). Similarly, Clark (2017) trained a neural network model to generate images in the style of google maps and satellite imagery from just an input map. Clark then used ancient city maps to feed the trained model and obtain interpretations of what those old maps would look like in a contemporary digital format.

3. Approach

Our system generates stylised geographic maps (output maps) with control over the location, visible features, and style. Because the system uses neural style transfer to generate the output maps, it can be split into two primary inputs: the input map and the style image. The input map is generated inside the system, requiring only a geographic area composed of two sets of latitude and longitude coordinates — the top-left and bottom-right corners of a rectangular area. Secondly, the style input is an image file not limited by pixel resolution. Granted that this image controls the overall look of the output thus, whilst the system accepts any image, not all of them will produce good results. The output maps are generated in a lossless raster image format (PNG) at pre-adjustable pixel dimensions.

The following subsections provide a brief overview of the pipeline that turns the system's inputs into output maps. A diagram of the main components that constitute the system and the interactions between them is shown in figure 1. The components are numbered and will be referenced throughout the following paragraphs.

Fig. 1. Diagram showing the main components of the presented system.

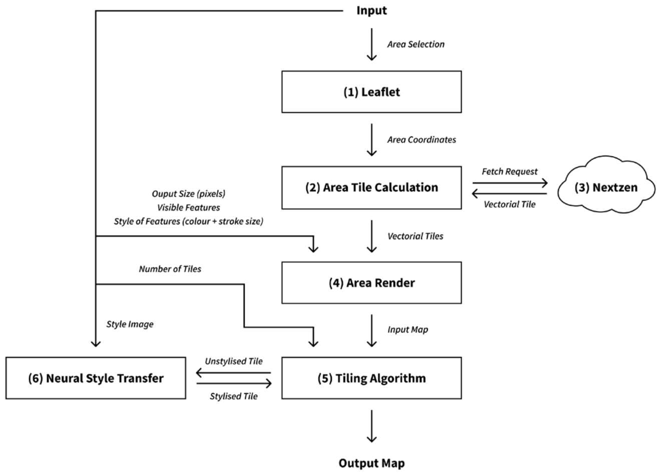

3.1. Rendering Input Maps

Our initial plan was to obtain input maps from the open-access data provided by OpenStreetMap (OSM), however, because we are not able to dynamically access OSM to retrieve area maps, we had to find a less direct solution. We followed a similar approach to the one used by most slippy map (i.e. interactive web maps that support panning and zooming) frameworks around the web, which is: get area maps through individual tile requests. In short, (in this context) tiles are square sections of a world map with varying zoom indexed by three values (x, y and zoom) and hosted on the web for frameworks to dynamically fetch, assemble, and display the section of the world map visible to the user at any given time.

To make area selection more intuitive and less bothersome, we used Leaflet[3] (component 1), an interactive map framework, to display an interactive world map with location labels through which we could visually select geographic areas. With the desired area selected, the system calculates which tiles compose that region and fetches (component 2) them from Nextzen[4] (component 3), a free service that hosts vectorial tiles based on OSM data and other open data projects. As opposed to bitmap tiles, where geographic information is compressed in an array of pixel values, vectorial tiles are not image files. Instead, vectorial tiles are geoJSON[5] objects that contain a list of geographic features with the corresponding latitude and longitude coordinates and indica-

3. https://leafletjs.com/

4. https://www.nextzen.org/

5. https://geojson.org

tors to correctly translate them to visual elements (points, lines and polygons). Considering these tiles are structured data, rendering a map from them is not as simple as arranging a few images in the correct order. Because of this, we implemented a tile rendering algorithm (component 4) that uses the pseudo-Mercator projection[6] to take care of rendering the geographic areas selected using Leaflet. This algorithm can convert an however long array of geoJSON tiles into a single SVG map.

6. https://epsg.io/3857

In summary, our system can render vectorial SVG maps of any geographic area in the world. The use of vectorial tiles greatly influences the amount of customisation supported by the system because they allow us to: (i) filter geographic features, e.g. only show highways; (ii) colour and set stroke sizes, which influences how the neural style transfer model interprets the semantic content; and (iii) losslessly convert the input map into any pixel resolution.

3.2. Map Stylisation

With the process described in the previous subsection (3.1) completed, we should have a satisfactory input map with the desired geographic area, colours, stroke sizes, and visible features. At this stage, we can select a style image and the pixel dimensions for the output map and subsequently run the neural style transfer model (component 6) to generate an output map. The latter parameter (pixel dimensions of the output) is needed because said style transfer model requires its inputs to be bitmap images, therefore the input map is converted from the generated vector format (SVG) to a bitmap format (PNG). For the neural style transfer model we are using an existing implementation of arbitrary neural style transfer[7] by Reiichiro Nakano, implemented using TensorFlow.js[8] and based on a paper by Ghiasi et al. (2017).

7. https://github.com/reiinakano/arbitrary-image-stylization-tfjs/

8. https://www.tensorflow.org/js/

After running some initial tests, we noticed that, with our computer's specifications, the system was unable to generate output maps larger than roughly 1500x1500 pixels. Considering the input map is capable of representing a substantial geographic area with all its road, building, and path detail, 1500x1500 pixels is a small size. During those initial tests, we also noticed that the outputs are not stochastic — meaning that if the inputs are kept the same, input map X plus style Y will always produce the same exact output map — and, with this in mind, we devised an algorithm that uses a tile approach to allow the system to generate outputs larger than 1500x1500 pixels. Our tiling algorithm (component 5), diagrammed in figure 2, works as follows: (step 1) split the input map into an adjustable number of equally sized sections; (step 2) run each of them individually through the neural style transfer model; and (step 3) assemble the resulting tiles back into a map. Note that each tile is still constrained to a maximum of 1500x1500 pixels, but by increasing the number of tiles we can generate bigger maps (for example, a 6000x6000 pixels map can be generated from a 6x6 grid of 1000x1000 pixels tiles). As seen in figure 3, this approach proved viable but introduced some visual flaws in the borders between tiles.

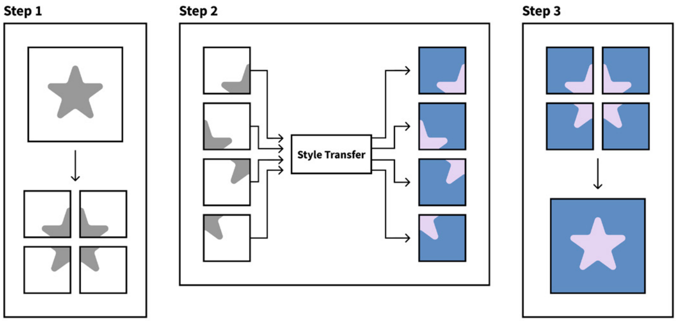

Fig. 2. Diagram visualising version 1 of the tiling algorithm used to generate outputs with larger pixel dimensions in the presented system.

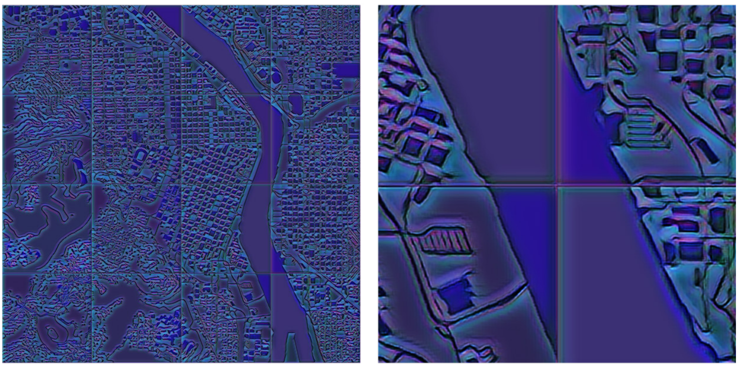

Fig. 3. Map of Portland, Oregon, generated with the presented system using version 1 of our tiling algorithm (left: entire output map; right: closeup section).

The lines visible in the tile borders are a result of the neural style transfer model, and we were unable to avoid them. As a result, we iterated over the tiling algorithm to solve the problem. The solution found (fig. 4) relied on generating each tile with a surrounding margin (step 1), which is cut out before the map gets assembled in the final step (step 3). This means that during the stage where visual flaws are introduced (step 2), each tile has a temporary buffer zone to catch them. These changes fixed the main tile border problem but, as shown in figure 5, revealed other (albeit less noticeable) visual inconsistencies between the tiles.

Fig. 4. Diagram visualising version 2 of the tiling algorithm used to generate outputs with larger pixel dimensions in the presented system.

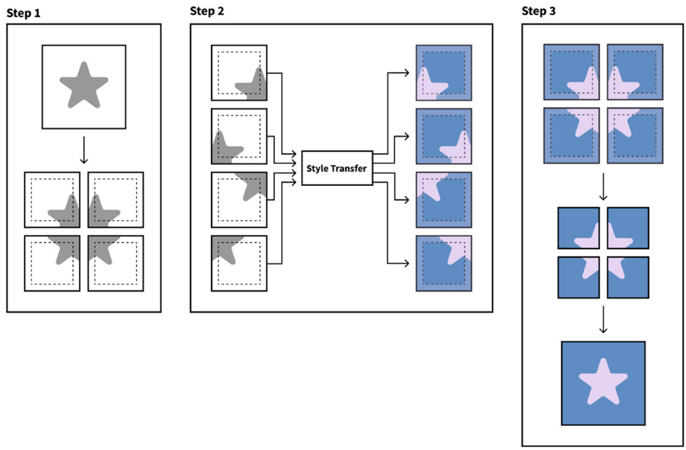

Fig. 5. Map of Portland, Oregon, generated with the presented system using version 2 of our tiling algorithm (left: entire output map; right: closeup section).

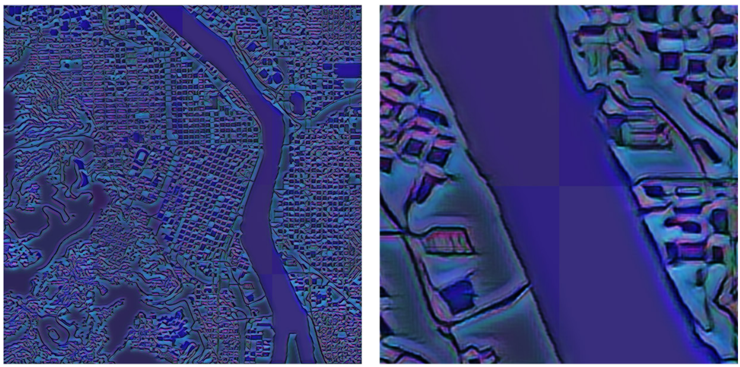

The remaining problems are most noticeable in the transition between tiles with varying amounts of negative space. Once again, these inconsistencies are caused by the neural style transfer model, and we were unable to avoid them, leading us to iterate over the tiling algorithm one last time. This time only the final step (step 3) was modified, where we faded between the tiles instead of slicing off the entire margin and leaving sharp edges between them (fig. 6). Finally, as visible in figure 7, this last version of the tiling algorithm seemed to fix the substantial inconsistencies around the transitions between tiles.

Fig. 6. Diagram visualising the final step of version 3 of the tiling algorithm used to generate outputs with larger pixel dimensions in the presented system.

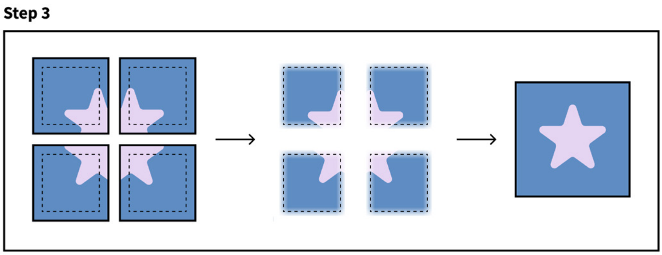

Fig. 7. Map of Portland, Oregn, generated with the presented system using version 3 of our tiling algorithm (left: entire output map; right: closeup section).

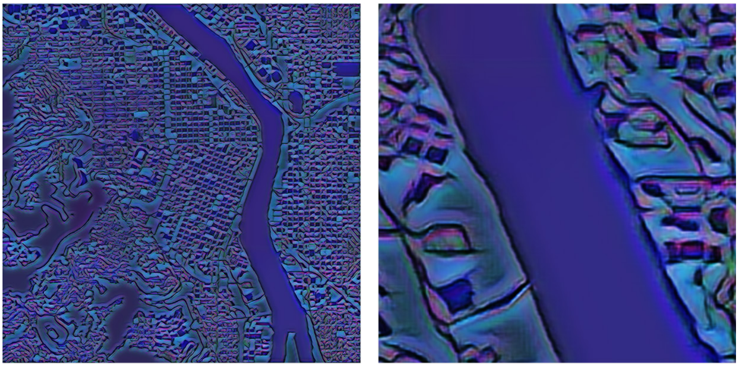

4.Results and Applications

With the system implemented, we ran some experiments to assess the visual quality of the output maps and how their appearance can be controlled by changes to the inputs. The following two subsections present the two main properties that control the visual quality of the outputs. For both subsections, the geographic area (fig. 8, left) used was kept the same because that is the one input that should not dictate the quality of the map (i.e. if we need a map of a specific city, it does not make sense to change the map location to improve the look of the output). Finally, the last subsection analyses the system and lists a few examples of how its output maps can be applied in real-world contexts.

Fig. 8. Map (left) and style input (right) used in the following examples. The map represents a section of Berlin, Germany, showing only buildings, roads and water bodies in black and white. The style is the illustration East Indian Cherry by Julie de Graag, 1919.

Map abstraction refers to how much the shapes that constitute the map differ from the original ones after stylisation. We found that abstraction is inversely proportional to the pixel dimensions of the neural style transfer output (larger images are progressively less abstracted). However, as previously explained, our system splits stylisation into a controllable number of tiles, which effectively means we have some control over the level of abstraction. Specifically, because abstraction directly correlates to the pixel dimensions of each image that passes through the style transfer model, we can increase the number of tiles which reduces the size of each one, resulting in higher levels of abstraction. For example, an output map of 2048x2048 pixels can be generated using a 32x32 grid of 64x64 pixels tiles (fig. 9, left) or using a 4x4 grid of 512x512 pixels tiles (fig. 9, middle), the output dimensions remain constant, but the level of abstraction decreases. We should emphasise, however, that we have limited control over this aspect. At a certain point, the number of tiles becomes so large that each one only represents a small fraction of the map, which introduces severe visual inconsistencies. For instance, if each tile is smaller than the average building, it becomes increasingly likely that there will be tiles completely filled with a solid colour, which the neural style transfer model interprets as being empty.

Fig. 9. Map of Berlin, Germany, generated with the presented system with three different levels of abstraction. From left to right, the geographic features are progressively more distinct, and the output is visually closer to the input map (shown in figure 8).

Abstraction affects the aesthetic quality and uniqueness of the outputs. Less abstracted maps show a clearer and more faithful view of space but, in doing so, lose what makes them unique. As seen in the rightmost map of figure 9, a map with low abstraction looks less like a stylistic rendering and more like a coloured version of the original one shown in figure 8. Customised maps that solely rely on colour and texture are already popular among the domain of map customisation and it is not our intention to further contribute to the over-saturation of these maps. We believe the best outputs to be those where the overall structure is easily recognisable but still retains a certain level of abstraction that makes it unique — in concrete terms, something like the middle image of figure 9. Due to the subjective nature of aesthetic preference, the level of abstraction we prefer is irrelevant, what matters is that our system has some control over that aspect in order to appeal to different aesthetic preferences.

4.2. Style Variation

After conducting a few experiments with different input styles, we concluded that (in our subjective opinion) the best maps stem from images with well-defined shapes, or better yet, images with borders (or strokes) between the various colours and shapes. Whenever the input possesses these properties, the neural style transfer model seems to be able to correctly translate the borders from the style to the output, resulting in most shapes having clear stroke delimitation. An apparent example of this can be seen in figure 9, where the dark bold borders of the style image (fig. 8, right) contributed to the resulting map shapes having dark edges — almost like the map was first traced with black ink and subsequently coloured following the black guiding lines.

That being said, bold strokes are not crucial to generate a clear, well-defined map. Effectively, the primary visual requirement of the style image is possessing distinct patterns and shapes. Such is the example shown in figure 10, where the thin golden strokes lining the shapes seem to be enough to generate a pleasing well-defined map. Conversely, images that lack this clear contrast and definition generate blurrier output maps with less distinct features. Unfortunately, this is the case for many hand-painted images where the natural hand-drawn brush strokes tend to leave a smooth transition between colours. An example of this is visible in figure 11, where the style has dark lines blocking the main shapes, but they fade somewhat smoothly to the surrounding colours, resulting in less clarity and shape delimitation in the output map. There are some areas where the map almost looks out of focus — especially when comparing both closeup sections of figures 10 and 11 — because colours dissolve into each other instead of having clear separations.

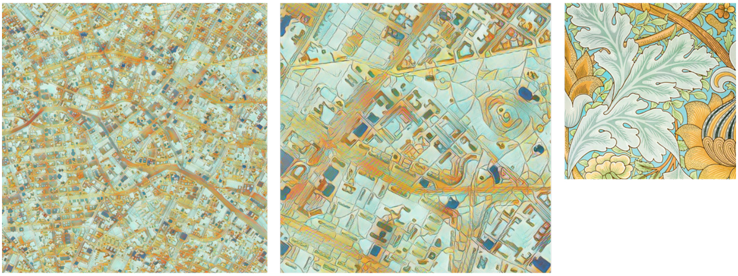

Fig. 10. Map of Berlin, Germany, generated with the presented system (left) with a closeup section (middle). The image used as the style input is the wallpaper St. James Pattern by William Morris, 1881 (right).

Fig. 11. Map of Berlin, Germany, generated with the presented system (left) with a closeup section (middle). The image used as the style input is the painting Pink Begonias by Marsden Hartley, 1887 (right).

4.3. Applications

Our first and primary motivation was to develop a system that generates visually appealing maps to be used as decorative objects. We feel this objective was achieved successfully as the maps (in our subjective opinion) do possess similar aesthetic properties to (albeit more abstract) paintings and posters commonly used to decorate interior spaces. In this regard, the output maps generated by our system could be used to decorate homes (fig. 12) and businesses alike. A real-world example of the latter would be their use to decorate the walls of a restaurant that specializes in traditional food from a particular city. Moreover, primarily decorative uses can be more than simple wall prints, there are a myriad of other contexts in which they can be used, of which souvenirs (e.g. postcards), city merchandising, and digital wallpapers are just a few examples.

Fig. 12. Map of Coimbra, Portugal, used for home decoration. The map was generated by the presented system and later digitally composed into a real-world scene using image editing software.

Ultimately, these output maps are a stylistic way of establishing a connection to a geographic space when there is no need for a readable map (i.e. when the aesthetic needs of the map supersede utilitarian ones). In this sense, our maps can also be applied in design objects that require a more or less subtle connection to a physical place, where it does not make sense (from an aesthetic or conceptual standpoint) to use a "traditional" map. Examples of these uses would be book covers, movie posters, or album covers, where the driving narrative is connected to a particular geographic place. To better illustrate the aforementioned design applications, we followed the first example and used our system to help design a mock-up cover for the 1988 book, The Twenty-Seventh City, by Jonathan Franzen. We generated a map of Saint Louis, Missouri (fig. 13), the city where the narrative takes place, with a style image that could help express the genre of the book — a complex thriller. Regarding the input map (fig. 13, left), we coloured the features so water and roads could be differentiated, thus allowing the bridges over the river to be visible in the output map. In the finished design (fig. 14), we cropped the output map and increased its contrast to better compose the book cover.

Fig. 13. Input map (left) and style image (middle) used to generate the output map (right) to be used in the book cover shown in figure 14.

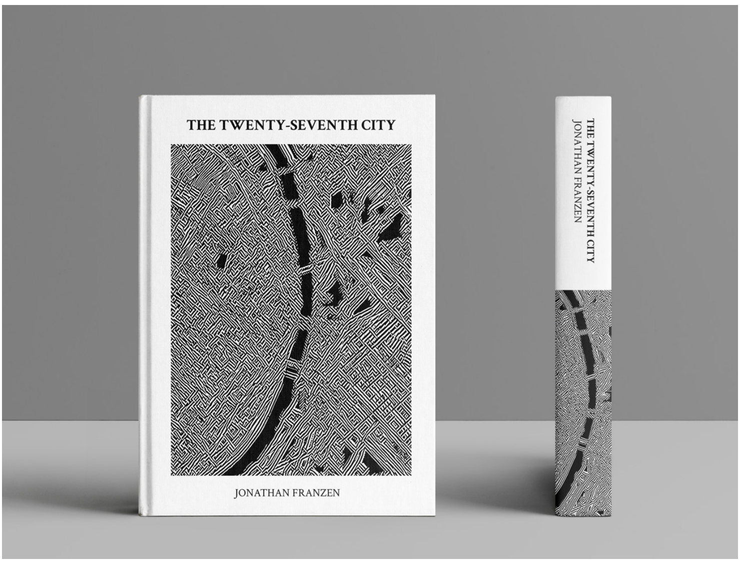

Fig. 14. Mock-up cover for the 1988 book The Twenty-Seventh City by Jonathan Franzen, using a stylised map of Saint Louis, Missouri, generated with the presented system.

5. Conclusions and Future Work

We have presented a system that uses neural style transfer to generate stylised geographic maps customised according to the style of a separate image. The presented system takes care of rendering the input maps with control over the visible features, their style, and the geographic area, thus requiring only a desired style image to generate stylised maps of any region in the world. Complete control over map area and location coupled with the freedom to use any image as the input style means the range of outputs is nearly limitless and highly customisable. We conclude that the generated maps possess visual interest and a unique aesthetic quality suited for both decorative and design applications.

Future work may include hosting the system on the web to allow anyone to create their own stylised maps. The system is already implemented in JavaScript and runs on the browser, therefore the main aspect that currently limits us from sharing the system with the public is the lack of an intuitive user interface. Additionally, should we choose to follow this approach, optimisations to the rendering process need to be implemented to decrease the time and computational power currently demanded by the system. Regarding added features, we believe the system would benefit from the possibility to add location markers to the map as a way of highlighting specific locations. We think this would help expand the output maps into more real-world contexts. For instance, a print like the one shown in figure 12 could be further connected to its owner by having markers or simple distortions in meaningful locations (e.g. the street where the person grew up). Future developments on this work will be available online at https://cdv.dei.uc.pt/stylised-maps/.

References

Bogucka, Edyta P., and Liqiu Meng.
2019. *Projecting Emotions from Artworks to Maps Using Neural Style Transfer*. Proceedings of the ICA 2 (July): 1–8.

Caquard, Sébastien, Barbara Piatti, and William Cartwright.
2009. *Special Issue on Art & Cartography*. The Cartographic Journal 46 (4): 289–91.

Clark, Jack.
2017. *Resurrecting Ancient Cities with CycleGAN*. Accessed February 7, 2021. https://jack-clark.net/2017/06/05/import-ai-issue-45/

Consilium Technology.
2018. *What If Google Maps Was Art?* Accessed February 7, 2021. https://consilium.technology/what-if-google-maps-was-art/

Gatys, Leon A., Alexander S. Ecker, and Matthias Bethge.
2015. *A Neural Algorithm of Artistic Style*. ArXiv:1508.06576 [Cs, q-Bio], September.

Ghiasi, Golnaz, Honglak Lee, Manjunath Kudlur, Vincent Dumoulin, and Jonathon Shlens.
2017. *Exploring the Structure of a Real-Time, Arbitrary Neural Artistic Stylization Network*. ArXiv:1705.06830 [Cs], August.

Kent, Alexander.
2017. *Cartographic Style and the Aesthetic Fix*. The Cartographic Journal 54 (1): 1–4.

Kogan, Gene, Gabriele Gambotto, Ambhika Samsen, Andrej Boleslavský, Michele Ferretti, Damiano Gui, and Fabian Frei.
2016. *Invisible Cities*. Accessed February 7, 2021. https://opendot.github.io/ml4a-invisible-cities/

Morris, Bill.
2016. *Deconstructed Planet*. Accessed February 7, 2021. https://billmorris.io/planet/.

Rees, Ronald.
1980. *Historical Links between Cartography and Art*. Geographical Review 70 (1): 60.

So, Chaehan.
2018. *A Pragmatic AI Approach to Creating Artistic Visual Variations by Neural Style Transfer*. ArXiv:1805.10852 [Cs], May.

Thrower, Norman J. W.
2008. *Maps and Civilization: Cartography in Culture and Society, Third Edition*. University of Chicago Press.