# Shan Shui in the World: A Generative Approach to Traditional Chinese Landscape Painting

Weili Shi\*

Parsons School of Design



Figure 1: Scroll of Blue-green Shan Shui in Manhattan, New York partially unfurled, together with a furled handscroll in a samite box. (Video documentary: https://vimeo.com/shiwl/shan-shui-in-the-world)

# ABSTRACT

Shan Shui in the World presents shanshui (山水, landscape) paintings of selected places in the world (Manhattan, New York in this first production) generated by a computational process based on geography-related information. This project revisits the ideas implicit in Chinese literati paintings of shan shui: the relationship between urban life and people's yearning for the nature, and between social responsibility and spiritual purity. For an audience living in an urban area, a traditional shanshui painting provides them with spiritual support through the depiction of the natural scene of elsewhere. With generative technology, however, Shan Shui in the World has the ability to represent any place in the world-including the city where the audience is-in the form of a shanshui painting based on geography-related information of the place. The notion that shan shui can exist right here not only underscores the contrast between the artificial world and nature, but also reminds the audience of an alternative approach to spiritual strength: instead of resorting to the shan shui of elsewhere, we may be able to obtain inner peace from the "shan shui" of our present location by looking inward.

**Keywords**: Chinese painting, generative art, non-photorealistic rendering, data visualization, cartography.

## 1 THE CHINESE IDEOLOGY OF SHAN SHUI

The concept of *Shan Shui in the World* derives from traditional shanshui paintings. The word shan shui  $(\coprod \pi)$ , literally translating as "mountain-water", is a Chinese abstraction of

natural landscapes. As an art form, shanshui painting<sup>1</sup> depicts shan shui using brush and ink. This tradition began during the fourth and fifth centuries, matured in Song dynasty (960-1279), and has been carried on through the following millennium. Various schools of technique and theory emerged in history, yet the essence of shanshui painting is clear since its early stage. Instead of striving to reproduce the landscape in exact details, shanshui painters distill the nature into semi-abstract, highly structuralized visual forms in order to express their state of mind in response to this idealized nature. Not only does the content of the painting possess spiritual meanings, but the brush-and-ink techniques, derived from that of Chinese calligraphy, also represent the virtue of the artist. The depiction of the landscape is done in scattered perspective-different parts of the same painting have their respective points of view. Therefore, the viewing experience of a large painting of shan shui, especially a long scroll, feels like a journey into the idealized natural scene.

Shanshui painting is a key element of what Chinese calls literati arts)—amateur arts of the scholars. Along with their education in liberal arts, the Chinese scholars were also trained in various forms of fine arts. As these people passed through the imperial civil examination, they became scholar-officials in the imperial bureaucracy, and at the same time, creators and audiences of art. In this way, art became an essential part of the spiritual life of the entire community of ancient Chinese intellectuals.

The mindset of these scholar-officials were shaped and influenced by two mainstream ideologies, back and forth. On the one hand, they were educated in the orthodox system of Confucianism, which emphasizes order and benevolence. Therefore, the scholar-officials had a strong sense of social responsibility, and a will to serve the country and bring peace to all people under heaven. On the other hand, however, they were also profoundly influenced by the naturalist ideas of Taoism. They believed that the ultimate truth lies in the ways of nature. They awed the nature, and had a mental intimacy with it. This feeling would have been even stronger when their Confucianist social ideals were defeated by the harsh realities of politics. At these

<sup>\*</sup> weili@newschool.edu

<sup>&</sup>lt;sup>1</sup> Note that in my writing, I use shan shui (two words) as a noun phrase and shanshui (one word) as an adjective.

times, they could only turn back to the purity of the nature for spiritual support.

As a result, the scholar-officials often felt a yearning for hermitage, a tendency to go back to the nature. But to the majority of them, the retreat could have only been a mental activity. The Confucianist Doctrine of the Mean admonishes decent people against acting in excess. Therefore, the scholar-officials would have ended up in their position, serving the society as much as they could in a moderate manner.

In this situation, the scholar-officials had to find alternative approaches to the nature while staying in the cities. They built naturalistic gardens to live in. They rebuilt natural scenes in a miniature scale into penjing. Or even a symbol of nature—a scholar's rock—would bring their minds closer to it. And of course, they found great pleasure and spiritual support from the creation and appreciation of shanshui paintings. The idealized natural scene not only symbolizes the eternity of nature, in the words of Michael Sullivan[1], but also reflects the virtues that the scholar-officials treasure. In fact, shan shui is the most popular content of traditional Chinese painting, and shanshui painting became the most sophisticated domain of Chinese painting, as a result of the spiritual pursuit of hundreds of generations of scholar-officials.

From the above reasoning, we can infer that the ideas implicit in shan shui are actually **the relationship between urban life and people's yearning for the nature, and between social responsibility and spiritual purity**. We have discussed in a way as if this relationship is a conflict. In a Chinese mind, however, harmony of these two sides is the proper goal to aim for, as the scholar-officials have done in their life and artistic practice. No matter which form it takes, I believe the urban-nature relationship is significant to not only ancient Chinese intellectuals, but also us urbanized modern people. Not only that, but the tension between the two sides of this relationship is even more intense in our time. While urban life has changed so much since the days of the scholar-officials, our spiritual approach should also be updated. There is an urgent need for us to find our contemporary "shan shui" in order to regain inner harmony.

Creative people have made their initial attempts to reinvigorate the old theme of shan shui for the modern context. Architect Ma Yansong proposed the manifesto of *Shanshui City*[2], envisioning our life in organic shaped buildings inspired by shanshui painting. Artist Yang Yongliang makes photo collages[3] which resemble shanshui paintings but are actually comprised of elements of the status quo of our urban life—a brilliant contrast between both sides of the urban-nature relationship. With the theme of shan shui regaining its audience, the ideology behind it is making its contribution to our contemporary world.

#### 2 GENERATING SHAN SHUI IN THE WORLD

My entry point to the art of shan shui is **the problem of elsewhere** in traditional shanshui paintings:

For an audience living in an urban area who resorts to shanshui paintings for spiritual support, this traditional art form always provides them with the depiction of the natural scene of somewhere else, because there is no shan shui in the cities for the artist to paint. As a result, in a sense, the comfort gained from the artwork can always feel like a misfire. According to this logic, the enjoyment of shanshui paintings can be understood as fleeing from the reality (or from social responsibility in the aforementioned relationship), because the mind is going elsewhere, where the depicted shan shui is.

To resolve this problem, I employ the contemporary approach of **generative art**. Through a computational process, artworks can be produced by an autonomous system based on an algorithm designed by the artist, with data sources used as raw materials. In



Figure 2: Height and area of the buildings in Manhattan, New York plotted according to their location.



Figure 3: Adjacent buildings merged into mountains, indicated by colors.



Figure 4: Outline of the mountains generated with data based on building information.



Figure 5: Mountains rendered in the style of ink-wash painting.



Figure 6: Mountains rendered in the style of blue-green shan shui.

order to remove the problem of elsewhere, the system can be designed to use localized geography-related information as its data source, and generate shanshui paintings according to that. Therefore, thanks to generative technology, *Shan Shui in the World* has the ability to represent any place in the world—including the city where the urban audience is—in the form of shanshui painting according to geographical data of the place.

This generative approach of art creation enables me to achieve what traditional shanshui painters could not match up to: depending on where this project is to be displayed, I can generate paintings of the shan shui in New York, or the shan shui in Beijing, or even the shan shui in Rome, where shan shui could not exist in our physical world. The generative paintings will be based on the data of the area of interest, and therefore have a certain level of objectivity. That is to say, in the generative parallel world of shan shui, shan shui can exist anywhere, and right here.

The notion that **shan shui can exist right here** is the decisive message delivered by *Shan Shui in the World*. It not only underscores the contrast between the artificial world and the nature, but also defuses it by revealing the generative relationship between the two—while people have built the artificial world based upon the nature, now we can also bring the nature back to the artificial world in some form. In the past, shanshui painting implies people's spiritual ideals in the depiction of natural scenes; from now on, generative technology broadens its object of representation into any kind of data in the world. Most importantly, *Shan Shui in the World* reminds the audience of an **alternative approach to spiritual strength**: instead of resorting to the shan shui of elsewhere, we may be able to obtain inner peace from the "shan shui" of our present location by looking inward.

# **3** THE GENERATIVE PROCESS

This and the following section introduce the creative process of the first production of *Shan Shui in the World*. In this production, the information of the buildings in Manhattan, New York was used to generate six shanshui paintings in two sets, following inkwash or blue-green painting styles, respectively. The generative images were printed and framed into traditional Chinese scroll paintings, and were signed and sealed in the classical way. Both the generative and scroll-making processes will be discussed. Note that while the text focuses on this specific production, it never the less represents the fundamental approach of the project even if the details change in future productions.

The way the generative shanshul paintings were created can be summarized in four steps: collecting the building data, merging the buildings into mountains, generating mountain outlines based on the building data, and rendering the mountains in the style of shanshul painting. The outcome of each step is indicated in Figure 2 through 6, with the last two figures showing two possible configurations of the rendering step.

The majority of the code for *Shan Shui in the World* was written in C++, with use of creative coding toolkit openFrameworks[4]. openFrameworks provides a unified API to functionalities such as HTTP request, data parsing, vector mathematics, and advanced computer graphics, which facilitates the implementation of the generative algorithm.

The building data was collected using Vector Tile Service[5] of Mapzen. It provides map data from sources like OpenStreetMap[6] organized in square tiles. Through its API, various kinds of geographical data can be downloaded in open formats such as JSON. An openFrameworks application was developed to download all the tiles covering Manhattan, and collect all the buildings within the boundary of the island while combining overlapping building parts, and save their ID,





(b)





coordinates, height, and area to disk. A total number of 47331 buildings in Manhattan were identified, as plotted in Figure 7(a).

With geography-related information collected, the next step is to transform it into elements of shanshui painting. In this production, the shanshui paintings only depict mountains that are generated by merging adjacent buildings. With this most essential element depicted, the other essential element—water—is implied by the negative space formed by it, as has been done in many classical shanshui paintings. As indicated in Figure 8, for every building, a merging radius was calculated from its area, and all buildings within this distance were merged with it into the same mountain. If one building should be merged with buildings from different mountains, then these mountains were merged into one. This process was implemented as an openFrameworks application. As indicated in Figure 7(b) and (c), with different configurations (merging radius and maximum mountain size), the same database of buildings can be merged into either a large number of various-sized mountains, or a small number of evensized mountains. Various merging configurations were needed by different creative intents: when generating the shan shui of a portion of Manhattan, the variance of the mountain sizes (and shapes generated in later steps) was emphasized, while for the panorama of the island, slightly fewer mountains were generated and more consistency among the mountains was attained.

The main generative engine of Shan Shui in the World was implemented as an openFrameworks application that accepts merged mountain data and renders it into shanshui paintings. At this stage, the generative mountains were only collections of buildings. In order to paint them, the outlines were needed. Instead of generating 3D models of the mountains and render them into 2D images, my approach is closer to painting. Although the process is 2D from the beginning, it still generates different views from different perspectives. As indicated in Figure 9, when a mountain is being looked at from different directions, the projection of its buildings onto the projection axis may change order. This order was followed when generating the mountain outline. As indicated in Figure 10, the top and bottom lines of the mountains were modulated by the height and area of its buildings, respectively, with a bell-shaped line as their basic shape, so that the modulated lines always have some resemblance to the shape of a mountain. In the generative engine, the outline data was stored in ofPolyline instances, and the built-in smoothing and simplifying functions of openFrameworks were used to fine-tune the outlines. Because of the nature of the geographical data, the merging process will always result in a small amount of tiny mountains which only contain a cluster of buildings that are far away from all the other ones. These mountains would probably get inelegant shapes from its minuscule data sets, and are not fit for a shanshui painting. I would examine the mountain shapes before rendering, and filter out the unsuitable ones.

The interface of the generative engine is shown in Figure 11, with the mountain outlines displayed. It can also display the buildings before and after merging, as well as the mountains rendered in two styles of shanshui painting. The resolution of the canvas is configurable. During runtime, there are multiple groups of parameters to control the generation of the mountain outlines and the rendering of the shanshui painting: I can pan, zoom, and rotate the map to get a desired composition; I can control the scale of the mountain shapes to (or not to) accommodates to the zooming level of the map; there are parameters to control the modulation and smoothing of mountain outlines, and tools to filter out undesired mountains; there are also parameters to control the rendering process. The generative engine is able to output its canvas contents as 16-bit TIFF images.

While the generative engine was implemented with openFrameworks in C++, the actual code that renders the shanshui painting was written in OpenGL Shading Language as fragment shaders, which run on the GPU and has the ability to post-process the image before it is outputted. With proper use of shaders, the flat and rigid mountain shapes generated by the openFrameworks code can be transformed into vivid and organic final results.

In order to facilitate sophisticated shading, ping-pong technique was employed in the rendering process. As indicated in Figure 12, in a multi-step rendering process, the source and result



Figure 8: Merging adjacent buildings into mountains according to their merging radii.



Figure 9: Building orders within a mountain according different camera directions.



Figure 10: Modulating the mountain outline.

frames buffers exchange places in every step, so that the intermediate result of the current step becomes the source of the next step. With this technique, multiple shaders can be applied to one image, and one shader can be applied multiple times. Several pieces of shader code were written for the generative engine, in order to transform flat shapes into gradients, add noise to either the lines or the shades, blur the image in different direction, and tint the image blue or green. The generative engine executes the shaders during runtime, and passing parameters to control them, so that multiple concise shaders can accomplish complex rendering tasks together.

The rendering of a mountain is summarized as follow: First, a white lining with the shape of the mountain is drawn to cover the mountains in the back. After that, the shade of the mountain is first drawn with a solid fill between the top and bottom lines, and is then replaced by gradients. Randomness and blurring are added



Figure 11: Interface of the generative engine in mountain outlines mode.

in order to emulate the ink-wash effect of Chinese painting, and multiple layers of noise are added to imitate the texture of brushwork. If the style of blue-green shan shui is chosen, the shade is tinted with a blue or green color calculated from the ID of the mountain, which is identical to the ID of the tallest building in it. Because neighboring buildings tend to have close ID numbers in the data source, the resulting blue-green distribution of the mountain colors also has some adjacency. The mountaintop line is then drawn with randomness in its shape, and variance in its width and brightness. Noise is added, and the line is blurred so that it blends well with the shade.

The mountains in a painting are rendered in a back-to-front sequence with alpha blending so that the front mountains cover the back ones correctly and blends with them naturally. Mountain opacity can also be applied according to their distance to the viewer, which adds aerial perspective effect to the image.

While shading is normally a real-time process, the above approach takes advantage of its rendering capabilities to such an extent that a significant amount of time is needed to process each frame. With a huge canvas (frame buffer) to paint on and thousands of mountains to paint, the rendering of a generative shanshui painting can take as long as an hour. Because of the size limit of frame buffer, the extra-long handscrolls were rendered in segments and stitched together in Photoshop.

Of the six generative shanshui paintings of Manhattan, New York, the two handscrolls portray the whole island in the same panoramic composition, revealing its rhythm in ink-wash (see Figure 13) or blue-green (see Figure 14) style, respectively; the four hanging scrolls depict two intriguing portions of the island in two traditional compositions (high distance, see Figure 15(a) and Figure 16(a); and level distance, see Figure 15(b) and Figure 16(b)) in both painting styles.

#### 4 THE CRAFT OF SCROLL-MAKING

The scroll format is deliberately chosen, not only because I consider the complete production cycle of physical (environment)-digital (generative images)-physical (scroll paintings) meaningful, but also because of the beauty of this



Figure 12: Ping-pong technique of using two frame buffers as source and result alternatively.

traditional carrier of art. A hanging scroll is meant to be hung on a wall and viewed from a distance, while the the enjoyment of a handscroll requires the audience to unfurl it one segment a time— a much more intimate experience, as shown in Figure 1.

The blue-green paintings were printed on silk, while the inkwash ones were printed on rice paper—both are traditional materials for Chinese painting. As can be seen in Figure 17, the physical material gives the image an additional layer of natural texture. The prints were then framed into scrolls. These steps were done in a workshop in Beijing.

A scroll painting not only contains the painting itself, but also has inscriptions and seal imprints on it. These can be done by the artist themselves, the owners of the artwork across time, and even connoisseurs whose words affirm the value of the work. Writing the inscription requires training in Chinese calligraphy, while a seal can be carved either by its user or by someone else. Figure 18 shows two seals carved from stones, their imprints, together with a seal ink paste and a carving knife. Both calligraphy and seal carving belong to literati arts. I inscribed each scroll with the title of the work, my name, and the year of creation. For the handscrolls, the name of the project (Shan Shui in the World) was also written in a larger size. I carved two seals with my name and the name of the project, and imprinted them on the scrolls.



Figure 14: Scroll of Blue-green Shan Shui in Manhattan, New York. 2016. Handscroll. Ink and colors on silk. (178 x 12 inch)



(a)

Figure 15: (a) Downtown Manhattan, New York, High Distance. 2016. Hanging scroll. Ink on paper. (24 × 55 inch) (b) Uptown Manhattan, New York, Level Distance. 2016. Hanging scroll. Ink on paper. (24 × 55 inch)

## 5 DISCUSSION AND FUTURE WORK

It is interesting and revealing to look at the generative shanshui paintings of Manhattan and compare them with the actual geography. Thanks to its high density of buildings, the shape of Manhattan is easily recognizable in the panoramas. Large buildings tend to merge into big mountains, and smaller buildings are represented as hilly areas. While the artificial part of Manhattan (its buildings) has been transformed into generative nature (the mountains), the natural part of the island (Central Park) is interestingly left blank in the paintings. The rhythm of the city (skyscrapers in the Financial District and midtown, smaller buildings in other districts, and further enhanced by Central Park) contributes to the lively composition of the generative images. Therefore, to an audience who understands that what shanshui painters try to depict is always the spirits of the scene, the authenticity of the generative shanshui paintings of New York should be well-founded.

Although being recently generated by an algorithm, these shanshui paintings do carry some essential qualities of this classical art form, and may provoke the audience into a certain feeling of nostalgia. I hope that this project would provide the



Figure 16: (a) Blue-green Downtown Manhattan, New York, High Distance. 2016. Hanging scroll. Ink and colors on silk.  $(24 \times 55 \text{ inch})$ 

(b) Blue-green Uptown Manhattan, New York, Level Distance. 2016. Hanging scroll. Ink and colors on silk. (24 × 55 inch)

audience with a new way to look at our artificial living environment, and that the messages delivered by it could empower people in their spiritual path.

I view Shan Shui in the World as a long-term project with various possibilities. I am envisioning the development of the project in different aspects, and am looking forward to opportunities to carry it out.

First of all, generating shan shui in other places. The project is conceptualized as being able to generate shanshui paintings for any place. It is natural to generate shanshui paintings for the place where the project is to be displayed. Also, like traditional shanshui paintings, the composition of the generative paintings does not need to literally follow the geography of the interested area. It would be reasonable and interesting to have a handscroll of shan shui along a river, as many traditional painters do, or even of shan shui along a subway line, or a highway-the generative shan shui does not have to be based on buildings of cities. Although it might be criticized as over-simplification, distilling the geographical information of a city into the height and area of its buildings is effective and expressive, because it matches up well to the



Figure 17: Details of *Downtown Manhattan, New York, High Distance*, featuring the inscription and seal imprints.

audience's impression and expectation of the city. When introducing new data sources, the choice need be made with a good understanding of the place, so that the data and its generative representation have a meaningful connection. For example, for the shan shui of San Francisco, it would make sense to include ground elevation along with building heights when generating mountain outlines.

Second, refining the rendering techniques. The current production already represents the essential aesthetics of shanshui painting. In future productions, the brush-and-ink techniques of Chinese painting could be explored in more depth. For now, a great part of interest of the generative paintings comes from the rhythm among multiple mountains. It could also be explored how to generate richer details for a single mountain. In the current paintings, only mountains were explicitly depicted. More elements, such as water, clouds, trees, boats, houses, and people could also be generated from various data sources. Technically speaking, the non-photorealistic rendering (NPR) process of Shan Shui in the World can be improved in many ways. While computational simulation of Chinese painting has not been widely practiced, existing work on the rendering of watercolor [7] is nonetheless valuable reference for my further exploration. An alternative but promising approach is using machine learning algorithms to analyze and imitate the style of the great masters in history [8]. I believe that for an art project, technological sophistication is not the primary consideration, and should always serve the artistic goal of the project. Nonetheless, generative shanshui painting does not have to mimic traditional shanshui painting in every aspect. Like every art form, it gradually develop its unique aesthetics based on (and be true to) its own creative approach.

Last but not least, **alternative forms of presentation**. While the current form of traditional scroll painting works well, I am open to various physical and digital forms of presentation. Instead of still prints, it would also be great to display the generative shan shui in animation, or even exhibit live generative shan shui, especially projection mapping of the generative shan shui, especially projecting it back onto the buildings whose information the generative process is based on—would make the contrast between the artificial and the nature even stronger. It might also be helpful to provide a moderate extent of interactivity between the audience and the generative system, so that the concept behind the project could be better conveyed. The generative shan shui



Figure 18: Two seals used in the project, their imprints, together with a seal ink paste and a carving knife.

could go beyond the form of painting and into the third dimension—to be 3D printed into shanshui penjing (as introduced in Section 1).

# ACKNOWLEDGEMENTS

*Shan Shui in the World* is my thesis project in MFA Design and Technology, Parson School of Design. I would like to express my sincere gratitude to Marko Tandefelt and Loretta Wolozin for their wholehearted support as my thesis advisors.

The geographical data used by *Shan Shui in the World* is from © OpenStreetMap contributors, Who's On First, Natural Earth, and openstreetmapdata.com through Mapzen.

## REFERENCES

- [1] M. Sullivan, Symbols of Eternity: The Art of Landscape Painting in China. Stanford: Stanford University Press, 1979.
- [2] M. Yansong, Shanshui City. Zürich: Lars Müller Publishers, 2015.
- [3] Y. Yang, From The New World. 2014 [Online]. Available: http://www.yangyongliang.com/Photography/144.html?a=1
- [4] "openFrameworks." [Online]. Available: http://openframeworks.cc
- [5] Mapzen, "Vector Tile Service." [Online]. Available: https://mapzen.com/projects/vector-tiles/
- [6] "OpenStreetMap" [Online]. Available: http://www.openstreetmap.org
- [7] A. Bousseau, M. Kaplan, J. Thollot, and F. X. Sillion, "Interactive watercolor rendering with temporal coherence and abstraction," presented at the the 3rd international symposium, New York, New York, USA, 2006, pp. 141–149.
- [8] L. A. Gatys, A. S. Ecker, and M. Bethge, "A Neural Algorithm of Artistic Style," CoRR, vol. 1508, 2015.