Leveraging the Spatiotemporal Analysis of *Meisho-e* Landscapes

Konstantina Liagkou[1], John Pavlopoulos[2,3]*, and Ewa Machotka[3]

[1] Uni Systems, Greece konstliagkou@gmail.com
[2] Athens University of Economics and Business, Greece
[3] Stockholm University, Sweden
ioannis@dsv.su.se, ewa.machotka@su.se

Abstract. Japanese early-modern woodblock prints depicting pastoral views of the countryside, so-called *meisho-e* (images of famous places), are often defined today as landscapes (*fūkei*). However, the notion of *fūkei* is a modern cultural translation, which obscures specificities of Japanese visual culture, and intricacies of early modern spatiality or a socially produced space. To uncover these characteristics and provide a more nuanced understanding of *meisho-e* prints, we have engaged in a macroanalytical study of relationships between places depicted in prints and actual topography, aided by computational technologies rooted in Natural Language Processing (NLP). In our prior work, we experimented with automated harvesting of geospatial data from image-content-related inscriptions on two hundred prints. In this follow-up work, we undertake a large-scale automated mapping of *meisho* and we study the geographical distribution of sites featured in these prints. We explore two different computational paths, one using deep learning and one based on digital gazetteers, and reflect on the challenges and benefits of the applied computational approaches. We improve the former, which was the state-of-the-art, using pre-training, and we show that the latter is beneficial in terms of mapping. Finally, by using automatically extracted place-name entities, we undertake an analysis of prints over space and time. We release our code and the dataset for public use: https://github.com/Connalia/ai-jan-art

Keywords: NLP · Spatiotemporal Analysis · Art History.

* Corresponding author. The first author wrote the initial draft and conducted the experiments, while all other authors contributed equally through supervision and co-authorship.

1 Introduction

Among the most globally recognised artifacts kept and displayed in museums worldwide are Japanese landscape prints designed by some iconic artists such as Katsushika Hokusai (1760-1849) or Utagawa Hiroshige (1897-1858). These images depicting peaceful views of the countryside featuring mountains, rivers, and rural dwellings, are commonly assumed to depict realities and topographies of pre-modern Japan, and as such defined *fūkei* or 'landscapes'. However, *fūkei* is a modern term developed in the process of cultural translation shaped by Western modern art epistemologies. In fact, these prints are more appropriately described as *meisho-e* or 'images of famous places'. Initially, these 'famous places' (*meisho*) were not depicting actual sites, which could be geolocated on the map of Japan, but poetic rhetorical figures rooted in classical poetry. These so-called *utamakura* (lit. poem pillows) tied seasonal images and symbolic motives, with particular places [9].

In the context described above, the actual topography of *meisho* places was not the guiding principle for their visual depictions [4]. In fact, *meisho-e* prints such as those designed by Hokusai or Hiroshige curate geographical reality in multiple ways. They maintain links to this reality mainly through printed inscriptions that feature a wide variety of place-names. These characteristics changed in time, especially in the first half of the ninetieth century, when many new toponyms entered the world of printed culture, and *meisho* strengthened their relationships with physical reality. Nonetheless, identification of the depicted places and their geolocation is far from being straightforward, which challenges interpretation of *meisho-e* prints as landscapes, and hinders understanding of their social function in general.

The issues, which were only briefly presented above, have not been comprehensively addressed to date. The research on *meisho-e* prints remains fragmented and often focuses on specific print series or individual designers rather than attempting to look at the genre and its epistemology at large. This is mainly due to the richness and diversity of the visual material that escapes traditional analytical methods based on close reading or interpretation of selected individual images. In this context, Natural Language Processing (NLP) can facilitate the discovery of new knowledge, through the analysis of large cultural datasets of digitised objects, and offer a possibility to rectify this situation.

1.1 Research aims

This study provides a macro-analytical exploration of *meisho-e* prints through Named Entity Recognition (NER). The contributions of this work are as follows:

- We benchmark NER on inscriptions of *meisho-e* prints, comparing the state-of-the art [13], a fine-tuned Japanese BERT [6], with simpler gazetteer-based approaches. We show (i) that the former is better, and (ii) that further pre-training of BERT leads to better results, setting the **new state-of-the-art**.

- Despite its superior performance overall, we show that BERT-based approaches still fall short compared to the much simpler gazetteers, when geolocation of places is the objective.
- We extract place-name entities for approx. 20k *meisho-e* prints, which we use to perform **a large-scale spatiotemporal analysis** of place-names distribution across time, aiming to discover which *meisho* are the most popular per time period, and how these preferences were distributed in space; i.e., which areas were considered culturally significant and at which times.

As spatiotemporal analyses are not limited to the mapping of depicted places, we also conduct the analysis of formal aspects of prints aiming at developing an understanding of how space is represented in prints. More specifically, we take a pivotal step to study how the colour schemes in prints changed in time and in relation to the depiction of different types of places. This new experiment opens up new analytical venues that we plan to explore in the future.

2 Related work

Recent years saw an advancement in the field of Spatial Humanities and Spatial Art History building on Geographical Information Systems (GIS), Natural Language Processing (NLP) and Corpus Linguistics [18]. However, although these analytical paths bring good results in the study of contemporary datasets it is not always the same for historical materials. Also, spatial analysis of Japanese pre-modern materials such as *meisho-e* prints, remains especially challenging, among others due to their formal characteristics (e.g., different perspective principles), difficulties with place identification (e.g, ambiguity of the depicted visual motives, problems with accurate transcription and lack of textual metadata) etc. Important contributions have been made recently e.g. the digitization of Japanese prints collections and the development of print databases worldwide [1], and among others, the initiation of computational analysis of Japanese prints targeting the questions of style [17], attribution [8] and content e.g. images featuring figures of kabuki actors [20] etc. or even geolocating of the selected print series depicting Edo city [16]. However, these important efforts study relatively small datasets, focus on testing technical solutions rather than large theoretical questions and do not target 'landscape' images at a scale and depth that would enable the development of an entirely new epistemology of this art genre.

Large-scale spatial analysis and the mapping of *meisho-e* prints with geographical information systems remains challenging as place identification is facilitated by the image-content related inscriptions printed in the images that often feature place names. But due to the high complexity of the task, reading early-modern inscriptions has so far been conducted by experts in Japanese pre-modern art history and literature. The task has been challenging for both humans and machines and it is estimated that only 1% of pre-modern textual sources have been transcribed [14].

Computational tools are expected to improve this situation. Automated text recognition technologies, e.g., Optical Character Recognition (OCR) and Hand-written Text Recognition (HTR), can yield the text in an image in a machine readable form. However, conducting OCR and HTR analysis on pre-modern Japanese texts is a challenging task due to the intricacies of Japanese writing systems. There are several reasons for this. First, this is due to a large number of characters used in pre-modern texts (ca. 4,500 characters are used [12]), of which many appear only once or twice in a given dataset. Second, the texts use so-called *hentaigana* or 'kanji variations' in which kanji characters could be used alternately depending on their phonetic value, and in multiple ways or forms. As a result, the same word could be written in different characters [19] and in different forms, or they could be written in a phonetic alphabet. Third, methodological challenges also include a lack of training data and applicability of software developed for the analysis of Western materials and texts in the study of non-Western materials. Fourth, the layout of texts is not always sequential and is characterised by a great variety of spatial distribution across the page and integration with the illustrations (both in single prints and printed books).

Therefore, the application of OCR software for automated transcription of text has been developing slowly in the case of East Asian languages using logographic writing systems [3]. Only recently, the field noted considerable progress in the development of OCR instruments. The new tools for automated reading of early-modern texts such as KuLA (Kuzushiji Learning Application) [7], KuroNet and Miwo ([5]; [11]) and newly established databases, among others, the National Institute for Japanese Literature, the National Institute for Japanese Language and Linguistics, the National Diet Library Digital Collections, the Waseda University's database, the Ritsumeikan University's ARC Portal Databases as well as datasets created, among others, by the National Institute of Informatics and the Center for Open Data in the Humanities (CODH) provided an excellent incentive and facilitated progress in the field of automated data harvesting from Japanese pre-modern textual sources. However, it is not feasible to develop one computational transcription tool for all pre-modern manuscripts, printed texts in books and visual images.

These and other challenges are also relevant to the automated reading of place-names inscriptions on visual images, even the prints produced at the same period. Therefore, transcription of inscriptions is one of the main obstacles not only for historians and literary scholars but also for art historians interested in a large-scale spatiotemporal analysis of *meisho-e* prints. Therefore, to expand the existing analysis of prints, and achieve our research goals we engage the so-called 'distant viewing' approach [2] and explore technological solutions facilitating a large-scale automated mapping of *meisho* or famous places depicted in prints produced between ca. 1750 and 1850, and we study the geographical distribution of sites featured in these prints. We explore two different computational paths, one using deep learning and one based on digital gazetteers, and reflect on the challenges and benefits of the applied computational approaches.

3 The corpus of digitised *meisho-e* prints

The access to the data for this work was facilitated by the database hosted at the Art Research Centre at Ritsumeikan University, Kyoto. The Centre's digital databases of Japanese printed culture host approx. 700,000 (678,429) digitised objects kept at 28 institutions in Japan and abroad. We use the 200 train- and test-annotated data with their place-names from [13]. Our study used 22,959 inscriptions that related to keyword *meisho* (famous place). These digitised prints depict mainly natural environments. Only 10,421 of them have metadata specifying production dates important for our study. Statistical overview of textual information in the 200 inscriptions studied previously [13] as compared to our new set of 22,959 is presented in Table 1.

Table 1. Statistics of the inscription texts in our study and in [13], including the number of texts, the average, the min., the 1st and 3rd quartile, the max. length in characters.

	Num.	Avg	St.dev.	Min	25%	75%	Max
Ours	22,959	11	7.9	2	9	13	444
[13]	200	20	8.5	5	14	23	59

4 Methods

In this study, we used NER to conduct a large-scale automated mapping of *meisho* or famous places depicted in prints produced between ca. 1697 and 1978. BERT-based NER has been found to be accurate for the task in hand [13]. On the other hand, gazetteer-based approaches hold valid latitude and longitude coordinates while not requiring any labelled data, as is the case of deep learning.

BERT-FT In [13], the authors used a Japanese pretrained BERT.[4] The model was pre-trained on Japanese Wikipedia articles and was fine-tuned for place-name entity recognition. The authors showed that merging location and geopolitical entities into a single place-name one leads to improved inter-annotator agreement and results. We used this model as a strong baseline, geolocating the recognised entities on a map, in order to elaborate on the limitations of BERT-based models. We call this fine-tuned BERT as BERT-FT.

GOJ The Gazetteer of Japan (GOJ) issued by the Government of Japan,[5] includes more than 4,000 modern place-names. It does not include any historical place names. This is an important limitation, because the inscriptions of this

[4] `https://huggingface.co/cl-tohoku/bert-base-japanese`
[5] `https://www.gsi.go.jp/ENGLISH/pape_e300284.html`

study refer to such historical places, some of which may have been altered over time. Edo, for example, is now called Tokyo, and as such is found in GOJ.

GeoLOD In [10], the authors introduced a gazetteer of Japanese toponyms that comes bundled with a geotagging algorithm. The tool is based on data from three databases: "Prefectures of Japan", "Historical Administrative Area Data Set Beta Dictionary of Place Names", "Railroad Stations in Japan (2019)".

BERT-FP-FT The Wikipedia articles which were used to pre-train BERT-FT do not cover the language used in our inscriptions, which were produced during the Edo period (1600-1868). To address this issue, in this study, we used 20,346 inscriptions to further pre-train Japanese BERT, before fine-tuning for NER. We used a masked language modelling objective, which is a language modelling task where the model is trained to predict the missing token(s) in a text. Our hypothesis is that this objective will allow the model to learn the language and the context of inscriptions from the Edo period, used in our dataset.

5 Experiments

Masked language modelling, or MLM in short, was used to further pre-train the Japanese BERT. We used a batch size of 64, a max length of 128 tokens, a learning rate of 2e-5, 20 epochs, 500 warm-up steps, and a weight decay of 0.01. To measure the model's ability, we measured the accuracy of the predictions for masked tokens. Further pre-training improves the average negative log likelihood of masked tokens (the same for the two models), from 2.22 to 1.27 (-43%).

We used MLM to yield our BERT-FP-FT model, which we compared with GOJ, GeoLOD, and BERT-FT [13]. As shown in Table 2, GeoLOD was better than GOJ, achieving 39% in F1. Both gazetteer-based approaches achieved high precision but low recall. A preliminary error analysis revealed that they could not detect all place names from the Edo period, but more experiments are needed to verify this. BERT-FT achieved 77% in F1. BERT-FP-FT outperformed its competitors in all the evaluation metrics and its difference in F1 (+4) from BERT-FT, which we consider as the previous state of the art, is also robust, as we can see in Table 3, where we repeated the experiment three times.

	Precision	Recall	F1
GOJ	0.92	0.07	0.13
Geolod	0.97	0.25	0.39
BERT-FT	0.76	0.78	0.77
BERT-FP-FT	**0.79**	**0.82**	**0.81**

Table 2. Evaluation of GOJ, GeoLod, BERT-FT [13] and BERT-FP-FT.

	#1	#2	#3	AVG
BERT-FT	0.80	0.82	0.85	0.82
BERT-FP-FT	**0.82**	**0.85**	**0.90**	**0.86**

Table 3. F1 across the 3 folds used for Monte Carlo Cross Validation

6 Empirical analysis

6.1 Inscription text restoration

One of the major problems in the transcription of inscriptions on prints, which facilitates their spatiotemporal mapping, is the quality and readability of historical material. In time, the material quality of prints deteriorates e.g., as the result of light exposure the paper and pigments undergo discolouring, which hinders the readability of the texts. Prints also may be damaged in other ways (via tearing, insect activity), which diminishes their readability by the public (both experts and the wider public). In this context, MLM can help restore fragmented inscriptions, by replacing with [MASK] the token to be restored.

We present a use-case of this method by testing its applicability in the restoration of an inscription on a selected print by Utagawa Hiroshige Fig. 6.1. The print presents a view of the famous Seta Bridge in the southeast part of Lake Biwa with Mt. Mikami in the background. The upper left cartouche comprises the inscription (also shown in Fig. 6.1). We masked the second kanji character in this inscription assuming that it was not readable (e.g., destroyed) Fig. 6.1). We fed our MLM with "瀬[MASK]夕照" and the model correctly predicted that the missing character is 田 (Table 4, first two columns of the first row). By masking and restoring the first kanji character, however, the model had trouble identifying the correct character (Table 4, last two columns). This is probably due to the fact that the first character refers to the name of a specific place, while the second is a generic term 田 or 'rice field' often used in different place names in Japan.

瀬[mask]夕照		**[mask]田夕照**	
w/o	**w/**	**w/o**	**w/**
田 0.15	田 0.79	UNK 0.13	嶋 0.04
野 0.06	川 0.03	都 0.02	隠 0.03
下 0.05	崎 0.02	狩 0.01	茨 0.02

Table 4. Two masked tokens predicted by BERT w/o and w/further pre-training.

Similarly, in other prints featuring the inscription with the place-name 浅間山 or 'Mt. Asama' or 'Asama Mountain' (which is, in fact, a volcano), we observe that the word 山 or 'mountain' is correctly restored, but this is not the case when we mask the first word 浅間 or 'Asama' instead. This is because the word is a generic term describing a topographical formation (ie., mountain) used in

Fig. 1. The print by Utagawa Hiroshige (1797-1858) entitled “Sunset Glow at Seta” (瀬田夕照), from the series “The Eight Views of Ōmi” (近江八景), 1834-35, woodblock print, MET (OA). On the left we see two cartouches with inscriptions, the first one with the second kanji character masked, and the second one, the original version with no character masked.

many place-names (e.g. mountains) in Japan. Nonetheless, this experiment has shown both opportunities and challenges related to the automated transcription of inscriptions in *meisho-e* prints.

6.2 Data augmentation with OCR

The next analytical step towards large-scale mapping of places depicted in print is experimenting with computational tools with automated transcription of inscriptions on prints, which feature place-names. The corpus of Japanese premodern printed books and single prints is extremely rich. For example, Kokusho Sōmokuroku (General Catalog of National Books) alone includes more than 450,000 pre-modern books [15], 90% dated to the Edo period (1600-1868). However, only less than 1% of these books have been transcribed to date. OCR could provide a solution to this problem, presuming that historical books, prints, and documents are available in digital format. In this work, we hypothesise that NER can be performed on the OCR-recognised text, and information about the place-names can be extracted despite the noise generated during the automated recognition phase.

We used a pretrained Japanese OCR model,[6] in order to extract the text inscribed on *meisho-e* print images. Then, we applied NER on the OCRed output, investigating the possibility of enriching our primary source of data before moving on to exploration at a larger scale. For example, as is shown in Fig. 2, we analysed Hokusai’s print entitled “Inume Pass in Kai Province” (甲州犬目峠)

[6] https://huggingface.co/kha-white/manga-ocr-base

from the series "Thirty-six Views of Mount Fuji" (富嶽三十六景, ca. 1830-31). We extracted the white rectangular cartouche with the inscription, located in the upper left corner of the image, and we experimented with different OCR models to transcribe the text. Then, we applied NER on the OCRed text. The recognition tool performed relatively well, with the model recognising 富嶽 or "Fugaku" (marked in red colour) which is an alternative name for Mount Fuji. This is a promising result given that our model was not trained on OCRed input. If successful, this application will not only allow large-scale exploration, but can also unlock related applications for the study of early-modern printed books.

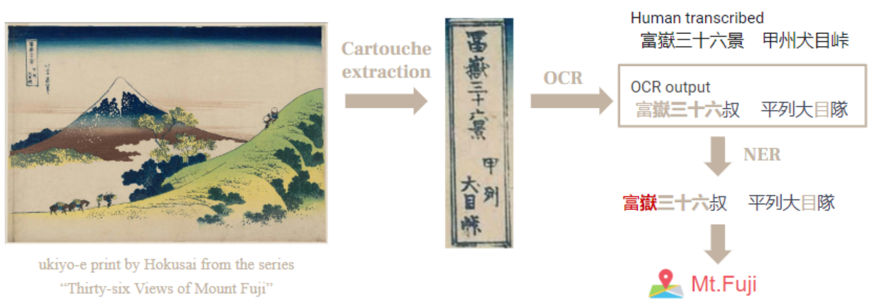

Fig. 2. Brown colour indicates the correct Japanese characters while red colour indicates the place-name extracted by the Bert NER model

NER on OCRed transcriptions is indeed promising, but challenges exist. When it comes to longer texts such as poems, due to the characteristics of the Japanese pre-modern writing system and the formal specificities of the design (e.g., multicolour cartouches), the recognition is largely distorted, for example, when the model delivered only four correct characters scattered across the inscription. The following question arises, then: is this error affecting the performance of NER models, and if so, can we learn to bypass it? To address this question, we plan to assess NER on OCR output, in order to quantify the error that is propagated from the recognition of written text to the recognition of named entities. Furthermore, we also plan to improve the written text recognition outcome. For example, one thing we are considering is joining forces with OCR error correction challenges.

6.3 Geolocating recognised place-name entities

Finally, we experimented with geolocating the recognised places, by applying several methods (as described earlier in Section 4 (Analytical Tools). The scope of this experiment was to assess if the reported evaluation results are reflected in a use case. In principle, the map featuring places depicted in prints, which would result from our spatiotemporal mapping experiments, could be envisioned as a tool for the study of spatial relationships between geographical places and their representation in prints and their changes across time. Fig. 3 provides such an example, by visualising the frequency of appearance of recognised place-names across our dataset.

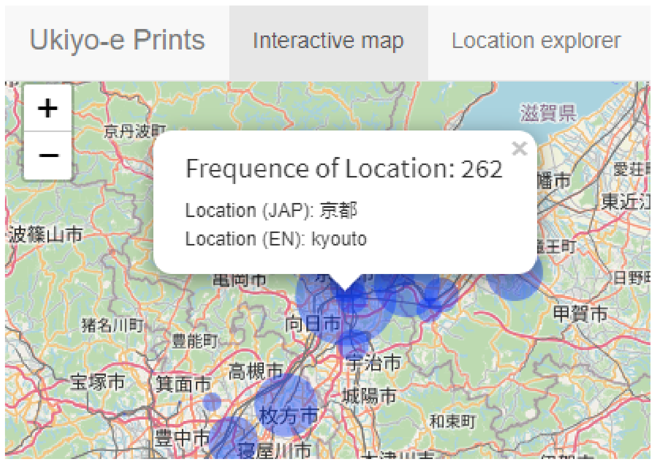

Fig. 3. Visualisation of a publicly available interactive map we created, serving exploratory purposes of the recognised (using GeoLOD) locations which *meisho-e* prints depict (the disc radius reflects the frequency)

BERT-based methods By using BERT-FP-FT, we observe that the model included different kinds of entities and not only place-names, yielding a faulty geolocation. In Table 5, which presents the ten most frequently recognised place-names across our data, we observe the presence of generic categories of landforms or human-made objects, e.g., 'mountain' or 'bridge', as well as grammatical forms such as the preposition 'of'. We also observe two more error sources. First, a few recognised places were geolocated outside of Japan, like Jiang, and Sichuan. Second, historical place-names extracted with NER from *meisho-e* inscriptions are not easily geolocated on contemporary maps of Japan. This is due to the historical transformation of Japanese writing systems, and historical changes in administrative geography in Japan as the retrieval algorithms are trained on contemporary datasets and gazetteers of toponyms (place names). For example, the tool could not geolocate the names of roads such as 'Tōkaidō' that are not pinnable on a map by a single pin. 'Tōkaidō' was represented as a larger grouping of pins and was located in Aichi prefecture. Also, Tokyo, Kawasaki, and Fujisawa were located on the Izu peninsula, which is not geographically correct.[7]

[7] We observed similar findings for BERT-FT

Table 5. The most frequent places extracted with BERT-FP-FT

Rank	Place	Count
0.	Tōkaidō	3657
1.	Edo	3328
2.	Toto	1602
3.	Tokyo	751
4.	Kiso	670
5.	of	545
6.	bridge	478
7.	Jiang	476
8.	Mountain	462
9.	Sichuan	452

Gazetteers When tested, gazetteer-based approaches, which already hold valid latitude and longitude coordinates, were overall more precise compared to BERT-based ones. The Gazetteer of Japan (GOJ) did not detect many place-names, mainly because it does not cover historical toponyms from the pre-modern period. GeoNLP was more precise but it could also not detect all place-names from the period. Therefore, despite the better performance of BERT-based models, we find that gazetter-based approaches still hold the advantage when the goal is geolocation and mapping of *meisho-e* prints.

6.4 Spatiotemporal analysis

Following our findings, we conducted spatiotemporal mapping of *meisho* depicted in prints aimed at the identification of the most popular places at different historical periods. In Figure 4, we present the frequency of recognised place-name entities over time. We used the GeoLOD for our purposes, to opt for high precision (Table 2). We have found out that the depiction of places (including both natural and man-made formations) flourished especially in the 19th century, especially between the 1830 and 1860s, and that different places were popular in different historical eras.

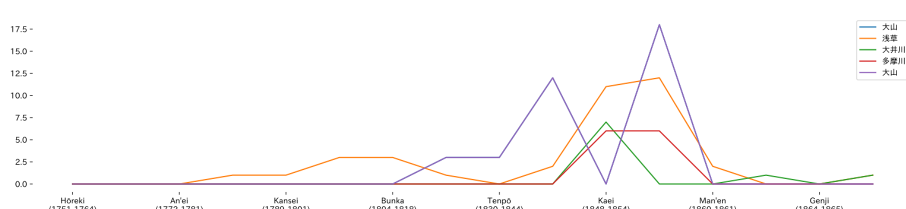

Fig. 4. Frequency of place name entities, recognised with GeoLOD, over time

We also experimented with analysis of formal aspects of the prints and how certain places were depicted, namely what colour schemes were used in their depiction across time. We analysed the colour of 3,505 *meisho-e* images printed during 16 historical eras, from 1751 to 1868. We transformed images from RGB (i.e., red, green, blue) to HSV (i.e., hue, saturation, lightness) and each pixel was classified to the colour it depicts, by using the classification presented in Table 6. This classification allows us to compute the percentage of each colour for an image (i.e., how many pixels are classified to the respective colour out of all in the image). Fig. 6 presents the colour percentages averaged across the images of each time period. We observe that yellow is the most frequently-used colour over time. Red and green follow, with the former being most dominating during early time periods (1751-1772) while the latter dominated after the Kaei era (after 1854). Blue and cyan followed, with low percentages over time. These analytical results may indicate historical changes in colour preferences among the print producers (designers, artists) and consumers as well as technological developments e.g. availability of certain pigments etc. They also may be correlated with certain types of motives and topographical formations (e.g., mountain, sea) and to a lesser degree even specific places (e.g., colour schemes linked to the depiction of the Edo city).

Table 6. HSV ([hue, saturation, value]) range (from, to) per colour

Colour	**HSV: from**	**HSV: to**
BLUE	[110, 50, 0]	[130, 255, 255]
CYAN	[80, 100, 100]	[100, 255, 255]
GREEN	[36, 0, 0]	[70, 255, 255]
RED	[0, 25, 0]	[15, 255, 255]
YELLOW	[16, 25, 0]	[35, 255, 255]

Fig. 6.4. For example, we observe that brown and blue dominate, which reflects the type of depicted motive, ie. the bridge, representing both water and land. Beige colouring also results in discolouration of the paper which prints undergo in time under the influence of the exposure to light. The dominance of brown and blue can also be seen in the 3D representation of the colours of that print (Fig. 6.4).

We investigated the correlation of specific colour schemes and historical periods and specific places. We focused on the two most frequent depicted places 東海道 or 'Tōkaidō' and 江戸 or 'Edo', which generated 468 and 575 hits respectively. We calculated the amount of blue and cyan per image and historicised the results Fig. 7. We observed that prints depicting stations along the 'Tōkaidō' road have more cyan colour than images of Edo city. This is not surprising considering that the 'Tōkaidō' road linking Kyoto and Edo city crosses large areas with rich waterscapes. Moreover, we observe that the colour is not directly linked to the place that is depicted, but rather follows the colour trends specific for a given era characteristic for the majority of prints produced at a given period.

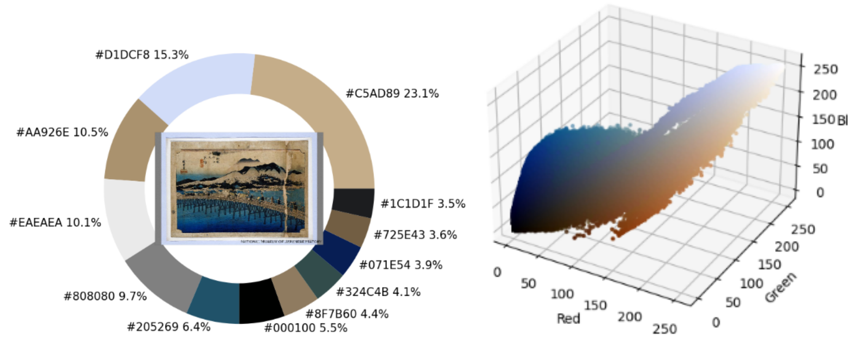

Fig. 5. Visualisation of the colour scheme in a selected print (title: 「東海道五十三次之内 京師 三条大橋」) focusing on the twelve most frequent colours (left) and a 3d RBG representation of all the colours in the image (right)

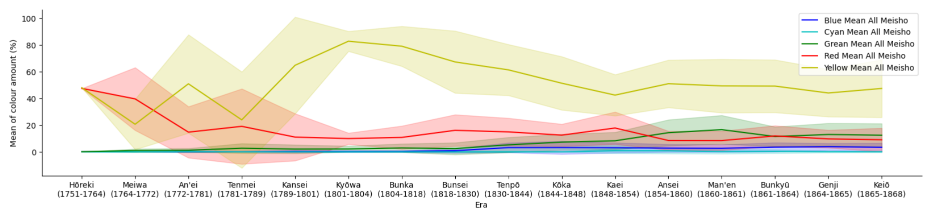

Fig. 6. Visualisation of the temporal distribution of the frequency of basic colours used in prints in different historical epochs between ca. 1750s and 1850s. with the standard error of the mean shadowing each respective time series

This finding underlines a general understanding of the prints and their material aspects in Japanese art history and indicates that to contribute to the state-of-the-art in the field in a meaningful more fine-tuned analysis needs to follow (e.g., on the level of a print designer, publisher, print format, topic).

Analytical limitations in the study of colour schemes The computational study of colour schemes in *meisho-e* prints has several analytical limitations. Most importantly, the colour schemes depend on the material characteristics and quality of the photographs of prints (subsequently digitised). Also, the results indicate a strong presence of silver and grey colour, which most probably is not included in the compositions but in the composition frames, which are also part of the image and its photo. These limitations need to be accounted for in the future interpretation.

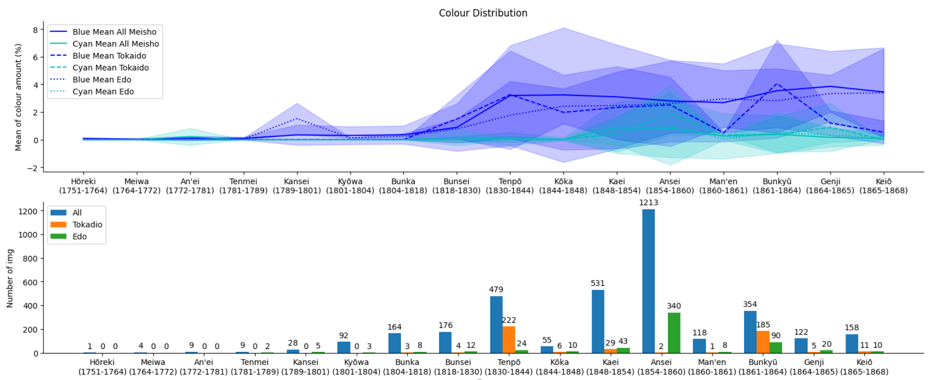

Fig. 7. On top is the average (and the st. error of the mean) percentage of blue and cyan in prints of specific time periods is shown, in the two most frequent places versus all images. The barchart shows the support (# prints) per place per era (i.e., aligned).

7 Conclusion

In this study, we benchmarked NER on inscriptions of *meisho-e* prints, comparing deep learning with simpler gazetteer-based approaches. We improved the former with further pre-training and we showed that BERT still falls short compared to the latter, when geolocation is the goal. By applying NER on approx. 20k images, we undertook a large-scale spatiotemporal analysis of recognised place-names, discovering popular *meisho* per historical era. By focusing on the two most frequently recognised *meisho*, we showed that colour reveals characteristics of the landscape, in the case of the waterscapes, and that it is not directly linked to the place depicted but may follow other trends. Future work will explore more *meisho* and we will interpret our analytical results, attempting to contextualise them in relation to sociopolitical changes and technological advancements (e.g. introduction of new printing materials), as well as correlate them with other factors such as print designer, publisher, and print format that may have played the role in their production and popularity at a given time.

References

1. Akama, R.: 赤間 亮, 立命館大学アート・リサーチセンターの古典籍デジタル化 : Arc国際モデルについて(¡特集¿古典籍資料の最前線), 情報の科学と技術. Journal of Information Science and Technology Association (2015)
2. Arnold, T., Tilton, L.: Distant viewing: Analyzing large visual corpora. Digital Scholarship in the Humanities 34:1: i3–i16. (2019)
3. Catalinac, A.: Quantitative text analysis with asian languages: Some problems and solutions. Polimetrics I (1) (2014)
4. Chino, K.: The emergence and development of famous place painting as a genre (2003)

5. Clanuwat, T., Bober-Irizar, M., Kitamoto, A., Lamb, A., Yamamoto, K., Ha, D.: Deep learning for classical japanese literature. Neural information processing systems (Neurips) creativity workshop (2018)
6. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
7. Hashimoto, Y., Iikura, Y., Hisada, Y., Kang, S., Arisawa, T., Kobayashi-Better, D.: The kuzushiji project: Developing a mobile learning application for reading early modern japanese texts. Digit. Humanit. Q. **11** (2016)
8. Hirose, S., Yoshimura, M., Hachimura, K., Akama, R.: Authorship identification of ukiyoe by using rakkan image. The Eighth IAPR International Workshop on Document Analysis System (2008)
9. Kamens, E.: Utamakura, allusion, and intertextuality in traditional japanese poetry (1997)
10. Kitamoto, A.: Toponym information platform geolod and its cooperation with rekiske and rekiroku (2022)
11. Lamb, A., Clanuwat, T., Kitamoto, A., KuroNet, A.: Regularized residual u-nets for end-to-end kuzushiji character recognition. SN COMPUT. SCI. 1, 177 (2020)
12. Lamb, A., Clanuwat, T., Kitamoto, A.: Kuronet: Regularized residual u-nets for end-to-end kuzushiji character recognition. SN Computer Science **1** (2020)
13. Liagkou, K., Pavlopoulos, J., Machotka, E.: A study of distant viewing of ukiyo-e prints. In: Proceedings of the Thirteenth Language Resources and Evaluation Conference. pp. 5879–5888 (2022)
14. Mitsutoshi, N., Susume, W.: 和本のすすめ : 江戸を読み解くために (2011), `https://cir.nii.ac.jp/crid/1130000795601158272`
15. Somokuroku, K.: Iwanami shoten **8 vols** (1963)
16. Suzuki, C., Kitamoto, A.: Pre- modern japanese books as data of humanities: Finding image of edo famous place from meisho- ki 名所記 and meisho- zue 名所図会 using iiif curation platform. Keynote Session (2019)
17. Tian, Y., Clanuwat, T., Suzuki, C., Kitamoto, A.: P2 - ukiyo-e analysis and creativity with attribute and geometry annotation. International Conference on Computational Creativity, Underline Science Inc (2021)
18. Won, M., Murrieta-Flores, P., Martins, B.: Ensemble named entity recognition (ner): Evaluating ner tools in the identification of place names in historical corpora (2018)
19. Yada, Tsutomu: 矢田勉. 国語文字•表記史の研究 (a historical study of japanese writing system). kyūko shoin (2012)
20. Yin, X., Xu, W., Akama, R., Tanaka, H.: A synthesis of 3-d kabuki face from ancient 2-d images using multilevel radial basis function. The Journal of the Society for Art and Science (2008)