TravelGANs.
A design approach to style transfer

Azzurra Pini [1]

Abstract. Generative AI has been used and researched mostly for artistic purposes and with a technology driven approach. We present an application of style transfer with GANs to design a service for the travel industry. Our solution is intended to support a mindset shift in the travel booking process, which is claiming for a more visual based search. TravelGANs is designed for sourcing local photos from the web and seasonalize them on the fly based on users' preferred travel dates.

1 INTRODUCTION

Generative AI has been around for quite a while to support human in specific tasks, mainly for data augmentation, to fill gaps of data and information in various contexts. Although the military has been a primary field of application for most AI techniques, the research has then evolved mostly in academia for computer vision research and applied to artistic projects.

Our project has been focusing in particular on the application of generative models, such as Generative Adversarial Networks, a class of machine learning systems invented in 2014 [6], that is now one of the most popular techniques used to learn patterns from data and reproduce its characteristics into new visual samples. Since then the technique has been used mostly to generate realistic photographs, text to image, video simulation, face ageing, image blending, super resolution and similar applications. At the technical level the capabilities of GANs, VAE and generative models in general have improved substantially over the years and we are now witnessing an impressive level of accuracy in reproducing realistic content. Nevertheless, only a few examples of real-world applications can be found, mostly for the fashion industry [3][13], for generating realistic models and apply different styles to clothing items. Some more rare ones are registered in the science domain [10] and engineering such as for autonomous driving, to simulate the possible different environmental context of roads in terms of lighting and weather. [1]

As opposed to the majority of the current research, instead of taking a technology driven approach, we tackled generative AI from a design perspective, looking for different ways where generative models can support a user experience or a mindset change.

TravelGANs tackles a behavioural change in the travel booking process which we claim to be more and more relying on visual features. The prototype aims at providing an AI powered visual booking service, capable of quickly sourcing local photos and seasonalise them on the fly in response to visitors' preferred travel period.

In this paper we will introduce the design concept, some experiments with the CycleGAN [14] style transfer and its implementation, the user interface and finally we will propose some directions for future research.

2 A-EYE SEASON TRAVEL

The project started observing a mindset change in travel booking towards a more visual experience, while the formats around communicating destinations have stayed static. Some recent initiatives by a few travel companies are have already acknowledged this cultural and behavioural shift, and started to move in that direction, providing for example users with the possibility to book travels through Instagram. [11][4][9]

With this idea in mind, the concept originated in particular from the idea of gathering travel inspiration from online pictures, and specifically to answer the question on 'how does (place) look like in (season)?'. Even the most visual platforms, more or less specifically dedicated to travel, do not in fact provide alternative photos of a place in different conditions – a functionality that we claim could be extremely relevant for travel planning, particularly for certain destinations with high seasonal variability.

From a technology perspective several options have been explored in order to tackle this challenge, ranging from computer vision techniques to more complex machine learning models. While classic computer vision techniques would support the analysis and editing of existing pictures, they are not accurate enough for transforming images without paired examples provided.

An interesting opportunity for experimentation has been identified in Generative Adversarial Networks (GANs), in particular we have been exploring the applications in relation to style transfer, an increasingly popular sub-set of algorithms originally started for artistic purposes [5] and now used to translate different types pictures into different styles, e.g. converting a photograph into the style of a famous painter or a horse into a zebra. More specifically, recent research has provided examples of application of GANs style transfer for translating photos of locations from winter to summer and vice versa. [14] Within this framework more recent research has also been made on transferring an image from night to day [1] and into different lighting or weather conditions. [1] [8]

The CycleGAN style transfer [14] is an approach for unpaired Image-to-Image translation developed at UC Berkley to learning to translate an image from a source domain to a target domain in the absence of paired examples. Among the different examples of

[1] Fjord, Accenture The Dock, Dublin, Ireland, email: azzurra.pini@fjordnet.com

transfer presented in the paper – the most popular being the artistic applications – we explored a later implementation into a season transfer for transforming pictures from winter to summer and vice versa. The current model is trained on 854 winter photos and 1273 summer photos of Yosemite downloaded from Flickr.

3 TravelGANs

With our prototype the challenge was to design a service to support a visual exploration of holiday destinations and for this purpose to exploit the potential of generative network. The service has been conceived mainly to support the travel exploration phase and it's meant to complement a regular booking platform (such as Booking.com). TravelGANs is meant to provide inspiration for holiday booking, in situations where there are less constraints in terms of destination and period with respect to other travel use cases. Specifically designed for ad-hoc holiday bookers, the service is meant to support an experience where users can find holiday destinations based on visual preferences rather than time and budget. The core concept is for them to understand how a certain destination will look like at the moment of the trip. The service is also tackling the use case of a B&B or small hotel owner as well as travel agents assembling packages. In this second case TravelGANs can help them rapidly assemble, customize and seasonalize their listings to increase the user engagement.

In order to test the concept we developed a web app prototype, using the current season transfer model of CycleGAN [14]. The design and development were performed in a rapid prototyping environment by a small team and delivered in a 6 weeks period.

3.1 Building a travel DB

From a software perspective we were aiming at the following high-level pipeline:

1. Selecting a set of popular travel destinations organized by type: seaside, mountain, cities, nature;
2. Building a static dataset for those destinations in different seasons, 3. Selecting winter and summer alternatives;
3. Run the images through the pre-trained model for season transfer;
4. Feed back the transformed images to the data base.

Our initial database has been populated by a web app we developed to access the FlickrAPIs, which allowed us to automate the collection of images, and in particular:

- Select a location
- Select a season
- Toggle for hemisphere (to support the seasons inversion)
- Filter licences for usage rights
- Pick preferred pictures and download a .zip file
- Feed the downloaded pictures to the CycleGAN season transfer model.

The web app is initially intended to support the development of the prototype, aiming at a future integration with a travel portal, but also to support the use case of the host, who can explore and select royalty free photos to add to his/her property page.

As displayed in Figure 1, based on the results of our test we also included some pre-processing steps that we found useful to improve

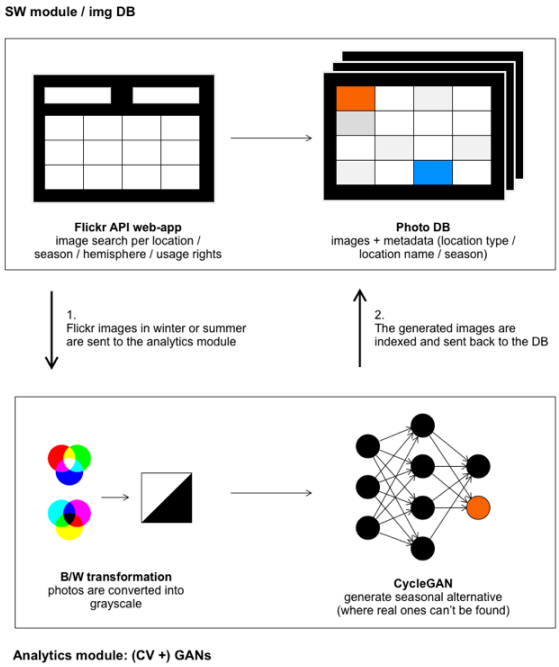

Figure 1. The implementation pipeline.

the transfer. The details of the season transfer tests and implementation will be discussed in the following section.

3.2 A CycleGAN application

The transformation of our pictures was made using the pretrained `summer2winter_yosemite` model[2].

Because the model has been trained mostly on pictures of Yosemite, when used to transfer seasonal style for other types of locations it visibly revealed a fundamental bias towards mountain landscapes of a "continental" latitude, with particular colour and lighting conditions. One of the challenges of our research was also to test the performance of the model on different types of landscapes.

Aware of this bias we also made some initial tests with a dataset containing a weather classification[3] to complement and improve the performance of the model, forcing associations between meteorological conditions and seasons. With the support of this classification we selected pictures tagged with peculiar seasonal weather conditions before feeding them to the model for the transfer.

Some tests returned particularly positive results [Figure 2], regardless of the similarity of the images with the training set, suggesting a valuable direction for future improvement of the season transfer. In order to gauge the applicability of the CycleGAN for the design, we tested the model also for other types of landscapes. Because of the limitations of the original training dataset we proved that the application of the transfer generally works better for pictures of mountains and similar natural landscapes. In particular we confirmed the hypothesis that the winter2summer transfer works pretty well with strong seasonal traits, e.g. snow, clouds, cold colours, thus supporting the utility of a complementary weather classification. At the same time, we discovered that generally the transfer doesn't work well when the original picture is very different from the examples of the training set, in the cases where the image doesn't present strong classic seasonal traits, i.e. if it's rather

[2] https://github.com/junyanz/CycleGAN.
[3] https://www.cs.ccu.edu.tw/~wtchu/projects/Weather/index.html

colourful and sunny. In this case sometimes the perceived effect is a reversed transfer, converting the image to winter (instead of summer). (Figure 3)

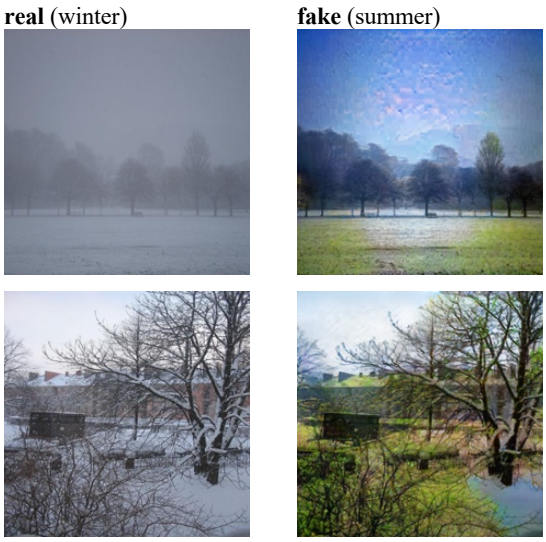

Figure 2. Examples of season transfer in combination with weather classification

Figure 3. winter2summer failure cases

A similar phenomenon unexpectedly happens also for environments that should have been similar to the training dataset, suggesting the need of a wider training also for these locations. (Figure 3b)
Since the prototype was designed to include different types of locations, further experiments have been made also with radically different landscapes. When an unexpected natural element is found (e.g. sand) the model fails to recognise it, and returns more "creative" results, mistaking for example a desert for sea and therefore replacing wrong colours (Figure 3c). As we expected the model doesn't also work very well with pictures of interiors, such as hotel rooms, that could be rather common for travel photos, especially as the pictures are gathered in an automatic way (e.g. through APIs).

Some of these experiments also revealed that the performance of the model increases significantly when the original picture is in black and white. In this case we hypothesize that the model is able to translate the images in a more straightforward way. Without the confusion created by unexpected colours, the model can simply apply the 'seasonal' style. Some tests of the winter2summer model on our dataset produced pretty good results. For this reason, we implemented the black and white conversion as a step in our pipeline before feeding the images to the model as displayed in Figure 1. As shown in Figure 4, after the conversion the hypothesised colourisation recognises and assigns pretty much all the right colours. Sometimes the fake images can be difficult to distinguish from the original colour version, but in terms of season transfer the results appear definitely more accurate and realistic.

Figure 4. winter2summer with black&white conversion

4 The User interface

TravelGANs' user interface is designed to support two different use cases. The main one we designed for is for a regular user who's exploring options for choosing a holiday destination.

The user is presented with a set of photos of travel destinations, randomly selected from our database through the Flickr API web app. The photos are arranged in an image board layout without any description indicating the location or season, to be used for visual inspiration. The user can refresh the selection and is prompted to select one that resonates with her/him.

After selecting a photo, the destination and the season of the image is revealed to the user. We use the metadata automatically associated to each photo - geographical coordinates, type of location and timestamp - to provide the user for other suggestions of similar destinations to explore.

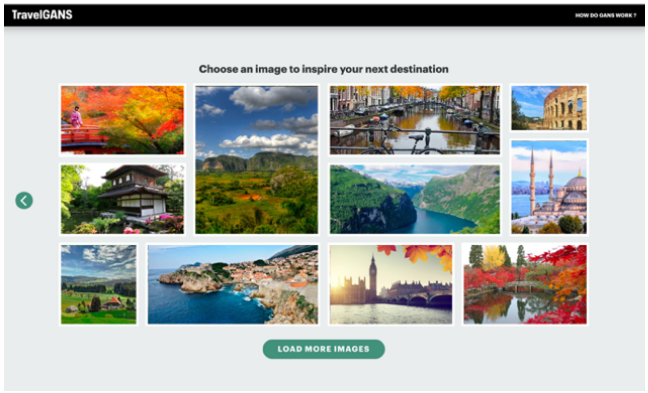

Figure 5. Photo board, destination selection

When a GAN generated image is selected, the original photo is displayed for comparison and the user is then guided towards the traditional booking process.

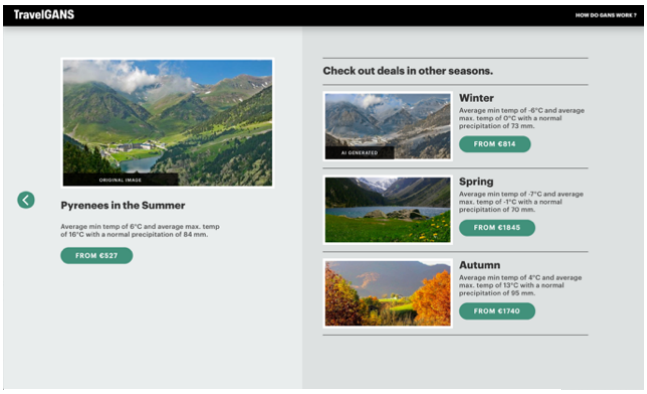

Figure 6. 'Seasonalization' of selected destination

5 FUTURE WORK

TravelGANs is a prototype of a service that suggests a new approach at holiday exploration, based on users' visual preferences. The implementation exploits the potential of deep generative networks to provide seasonal alternatives of travel pictures.

From a design perspective there is planned future work for an extensive user research, especially on the perception of AI generated images in relation to a travel booking service. Broader research is also required to design functionalities more focused on visual similarity and personalization of features.

From a technology perspective we discovered a great potential for further exploration with respect to GANs and more in general on deep learning models for generative design.

More specifically, we identified several opportunities for future work in:

- Training of different models for different location types, e.g. Seaside, cities and interiors.
- Training models on clustered locations based on similar visual characteristics, using external classifications such as the weather dataset [8];
- Implement computer vision techniques to analyse images upfront, such as image segmentation and classification integrating stuff classes [2];
- Contextual latitude/longitude time classification to improve seasonal accuracy;
- Semantic analysis and content recognition: recognising semantic elements in pictures would allow to apply different style rules to different objects (e.g. sea VS sand)
- Perform tests with similar generative techniques as Variational Auto-encoders (VAE).

The exploration of these technologies would allow also to expand the design of a more visual travel booking experience with complementary use cases to be explored.

- The simple application of content recognition would allow users for a semantic search of specific visual elements into travel photos. Such a travel service would therefore be able to respond to questions like "I want to go to a place with trees and museums".
- Dominant color extraction and similar techniques would allow to explore travel locations by their most frequent hues, responding to users' questions like "I want to go to a place that looks mostly yellow and blue".
- All these techniques can be integrated and provide live updated personalized visual cues and trends of the most searched terms or characteristics.

The project also became a starting point to deepen the reflection on human-machine interaction, in particular in relation to service design applications of generative AI and the need for parameters for the exploration and representation of the latent space.

REFERENCES

[1] Anoosheh, A., Sattler, T., Timofte, R., Pollefeys, M., & Gool, L. Van. (2019). Night-to-day image translation for retrieval-based localization. *Proceedings - IEEE International Conference on Robotics and Automation*, *2019-May*, 5958–5964.

[2] Caesar, H., Uijlings, J., & Ferrari, V. (2018). COCO-Stuff: Thing and Stuff Classes in Context. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1209–1218.

[3] Dietmar, J. "GANs and Deepfakes Could Revolutionize The Fashion Industry". *Forbes*.

[4] Coffey, E., Easyjet launches app allowing passengers to book flights based on Instagram pictures, (2018), *Indipendent.co.uk* https://www.independent.co.uk/travel/news-and-advice/easyjet-app-look-and-book-instagram-pictures-flights-social-media-a8588481.html

[5] Gatys, L., Ecker, A., & Bethge, M. (2016). A Neural Algorithm of Artistic Style. *Journal of Vision*, *16*(12), 326.

[6] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, Yoshua Bengio "Generative Adversarial Networks", in NIPS 2014

[7] Liu, M. Y., Breuel, T., & Kautz, J. (2017). Unsupervised image-to-image translation networks. In *Advances in neural information processing systems* (pp. 700-708).

[8] Lu, C., Lin, D., Jia, J., & Tang, C. K. (2017). Two-Class Weather Classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *39*(12), 2510–2524.

[9] Miller, C., How Instagram is changing travel https://www.nationalgeographic.com/travel/travel-interests/arts-and-culture/how-instagram-is-changing-travel/

[10] Mustafa, M., Bard, D., Bhimji, W. *et al.* CosmoGAN: creating high-fidelity weak lensing convergence maps using Generative Adversarial Networks. *Comput. Astrophys.* 6**,** 1 (2019).

[11] Ritschel, C., When booking holidays think of Instagram photos above all other factors, *Indipendent.co.uk*, (2017) https://www.independent.co.uk/life-style/millennials-holiday-decision-instagram-photos-factors-think-first-social-media-a8131731.html

[12] Chu, W. T., Zheng, X. Y., & Ding, D. S. (2017). Camera as weather sensor: Estimating weather information from single images. *Journal of Visual Communication and Image Representation*, *46*, 233–249.

[13] Wong, C. "The Rise of AI Supermodels". (2019) CDO Trends.

[14] Zhu, J. Y., Park, T., Isola, P., & Efros, A. A. (2017). Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. *Proceedings of the IEEE International Conference on Computer Vision*, *2017-Octob*, 2242–2251.