Toward XAI & Human Synergies to Explain the History of Art: The Smart Photobooth Project

Preprint · June 2021

7 authors, including:

Egberdina van der Peijl
University of Luxembourg

Amro Najjar
University of Luxembourg

Yazan Mualla
Université de Technologie de Belfort-Montbéliard

Yolanda Spinola-Elias
Universidad de Sevilla

Toward XAI & Human Synergies to Explain the History of Art: The Smart Photobooth Project

Egberdien van der Peijl[1], Amro Najjar[1], Yazan Mualla*[2], Thiago Jorge Bourscheid[1], Yolanda Spinola Elias[3], Daniel Karpati[1], and Sana Nouzri[1]

[1] AI-Robolab/ICR, University of Luxembourg, Computer Science and Communications, 4365 Esch-sur-Alzette, Luxembourg
egberdien.vanderpeijl@ext.uni.lu,thiago.bourscheid.001@student.uni.lu, {sana.nouzri,Amro.Najjar}@uni.lu

[2] CIAD, Univ. Bourgogne Franche-Comté, UTBM, 90010 Belfort, France
yazan.mualla@utbm.fr

[3] Dpto. Dibujo, University of Seville, C/Larana n°3, 41003, Spain
yspinola@us.es

Abstract. The advent of Artificial Intelligence (AI) has brought about significant changes in our daily lives with applications including industry, smart cities, agriculture, and telemedicine. Despite the successes of AI in other "less-technical" domains, human-AI synergies are required to ensure user engagement and provide interactive expert knowledge. This is notably the case of applications related to art since the appreciation and the comprehension of art is considered to be an exclusively human capacity. This paper discusses the potential human-AI synergies aiming at explaining the history of art and artistic style transfer. This work is done in the context of the "Smart Photobooth" a project which runs within the AI & Art pavilion. The latter is a satellite event of Esch2022 European Capital of Culture whose main aim is to reflect on AI and the future of art. The project is mainly an outreach and knowledge dissemination project, it uses a smart photo-booth, capable of automatically transforming the user's picture into a well-known artistic style (*e.g.,* impressionism), as an interactive approach to introduce the principles of the history of art to the open public and provide them with a simple explanation of different art painting styles. Whereas some of the cutting-edge AI algorithms can provide insights on what constitutes an artistic style on the visual level, the information provided by human experts is essential to explain the historical and political context in which the style emerged. To bridge this gap, this paper explores Human-AI synergies in which the explanation generated by the eXplainable AI (XAI) mechanism is coupled with insights from the human expert to provide explanations for school students as well as a wider audience. Open issues and challenges are also identified and discussed.

Keywords: AI & Art, XAI, agents, Neural Style Transfer, Cultural Heritage

* Corresponding author.

1 Introduction

In the last decade, AI has become omnipresent with applications spanning from autonomous vehicles, agriculture, and industry. In recent years, this wave has also spread to new domains such as digital history and cultural heritage. In this particular context, making the cultural heritage more accessible and more engaging by bringing it closer to more audiences is considered to be one of the main contributions of the latest generation of AI systems [15]. For instance, MonuMAI is a smartphone app allowing for artistic knowledge dissemination. MonuMAI classifies photos (*e.g.,* taken for a monument's facade) and classifies it into different architectonic style and provide visual hints explaining why the photo belongs to the style. To do that, MonumMAI relies on deep learning classifiers supported by expert knowledge [36]. Other works in the literature aim to increase the accessibility of cultural heritage for target user groups. This can be either for users with special needs (*e.g.,* elderly people with mobility constraints [33, 34]), or to enhance the level of audience engagement by providing comprehensible and interactive content. In this paper, we present the Smart Photobooth project, a project within the context of Esch2022 European capital of culture.

The Smart Photobooth is an outreach project aiming to disseminate knowledge about the history of art and artistic styles as well as the latest machine learning mechanisms and their applications in the domain of art generation and classification. The smart booth relies on Neural Style Transfer (NST) [30] to transfer an input image (typically the portrait of the user) to one of the most famous artistic styles (*e.g.,* cubism). NST relies on machine learning Generative Adversarial Networks (GAN) to achieve the transfer [30].

After the user gets their transformed portrait, they receive both *(i)* a short tutorial explaining the style and its position in the history of art based on well-known paintings from the style and *(ii)* an AI explanation highlighting the style features appearing in the transferred user portrait.

Recently, eXplainable Artificial Intelligence (XAI) has been identified as a powerful approach to support these outreach efforts by helping to interpret the otherwise incomprehensible inner-workings of sophisticated machine learning mechanisms such as the advanced GANs powering the NST process [15, 12]. In particular, XAI has been suggested as a possible solution to "explain a given artwork's success in terms of the underlying influencing artistic styles" [15].

However, this is a challenging task since, in contrast to many other disciplines (math, physics, chemistry), making art "understandable" requires a combination of objective and subjective interpretations to analyze its message. For instance, analyzing the objective features of Guernica, the famous painting by the Spanish artist Pablo Picasso is not enough to understand the background and the interpretations of this painting. The latter information is highly subjective, controversial, and depends on the historical circumstances surrounding the creation of the painting [48]: the Spanish Civil war, Picasso being commissioned by Manuel Azana, the president of the short-lived Spanish Republic, to create a large mural for the Spanish pavilion at the 1937 Paris World's Fair, as well as the bombing of Guernica, a town in the Basque country, on the 26th of April, 1937 by the Condor Legion of the Nazi German air forces, and Picasso's discussion with his friend the poet Juan Larrea who urged him to make the bombing his subject [38], *etc.*

To overcome this challenge, in this paper we propose a human-agent architecture allowing us to foster the needed synergies between the involved parties. Namely, the human end-user, the artist, and the black-box machine learning mechanism. Based on the context, the user preferences, and the artists' recommendations, the agent provides personalized explanations combining machine learning interpretation, agent explainability, as well as the artist's expert analysis. The architecture is discussed, the challenges it raises and are identified and discussed.

The rest of this paper is organized as follows. Section 2 lays out the background for this work. Section 3 introduces the Smart Photobooth project. Section 4 presents the proposed architecture, Section 5 identifies the challenges and the open issues, and Section 6 concludes this article.

2 Background

2.1 AI & Art

Works combining AI and art in the literature fall into two categories. The first is using AI in the process of creating new art while the second is using AI to analyze existing human-created art.

AI Art Generation The recent rapid evolution of Deep Neural Networks (DNN) has accelerated the use of AI technologies to create art. In particular, GANs are among the cutting-edge technologies used in this domain. The latter involves a couple of systems of DNNs designed to compete against each other. For instance, in the case of GANs generating visual art, a DNN, called the generator, is trained to generate realistic images whereas the other DNN (known as the discriminator) is trained to classify generated images as fake while identifying real artistic pieces as real art. The training of the GAN is achieved once the generator becomes capable of creating output that cannot be identified as fake by the discriminator (*i.e.,* the generator outperforms the discriminator). Recently, this type of models has been implemented in different configurations (*e.g.,* CycleGAN [60], StyleGAN [31], BigGAN [7]) and has achieved remarkable results in generating human faces (2D [45] and 3D [55]), music [35], as well as furniture [58]. Another application of AI Art generation is transferring some features of the input. Face aging and NST are notable examples.

Understanding AI with Art The advent of AI has a significant impact on art access ability and understandability. Several works in the literature have explored how AI can help make cultural heritage more accessible for users with special needs. For instance, haptic interfaces have been proposed to be used by museums to help visitors with visual impairments formulate mental pictures of the objects and provide important contextual and navigation information [56, 10, 13, 18] (*cf.* [47] for a review).

Moreover, in recent decades, thousands of artworks have been digitized and are now available online for access and analysis. OmniArt [53] and Art500 [40] are among the biggest artwork datasets. The former includes about two million artworks allowing for author, style, period, type and iconography retrieval, color classification, and object detection, whereas the latter contains about 550 thousand artworks. The Metropolitan Museum of Art of New York also released in 2017 over 406.000 indexed pictures of public domain artwork [1]. The datasets can retrieve items by authors, genre, styles, events, and historical figures.

Artistic Style Classification Based on the datasets mentioned above, recent works in the literature propose to classify paintings into their artistic styles (Renaissance, Baroque, Impressionism, *etc.*). Many of these works rely on style patterns and definitions proposals by the Swiss art historian Heinrich Wölfflin [57]. In particular, Wölfflin identifies five key visual principles each defined by two contrasting visual schemes [57, 11]:

(i) **Linear *vs.* Painterly**. In the former, elements are clearly outlined and boundaries are clear while in the latter, elements are fused and contours and edges are blurry. *(ii)* **Closed *vs.* Open forms**. In the former, elements are balanced with the frame. Vertical and horizontal compositions are dominant. In the latter, diagonal components are dominants with an impression of the space going beyond the edges of the picture. *(iii)* **Planner** in which elements are organized in successive planes parallel to the picture plane versus **Recessional** which gives an illusion of depth and where elements are arranged on various planes. *(iv)* **Multiplicity.** Elements appear distinct and independent, versus **Unity** where elements are fused into a single whole. *(v)* **Absolute clarity** with explicit and articulated forms versus **relative clarity** with less clearly structured forms avoiding objective clearness in an intended manner. Recent works in the literature suggested that convolutional neural networks trained to classify paintings according to their artistic style, implicitly learn features related to Wölfflin's concepts. For instance, Elgammal et al. [17] used a convolutional neural network to classify paintings into their artistic styles. The results obtained showed that the network managed to smoothly place an artwork into a temporal arrangement based on learning style labels. In a related work, the authors in [11] trained a convolutional neural network to predict the values of the five Wölfflin concepts (or features). The result of this work showed that the proposed network learned to discriminate meaningful features corresponding to the visual characteristics of Wölfflin's concepts.

But, according to other authors, such as Lecoutre, Negrevergne and Yger [39] or Tan et al. [54], identifying the artistic style of a picture in a fully automatic way is a challenging problem since classifying visual styles cannot rely on any definitive feature. This is especially difficult for non-representational artwork.

For these reasons, in this work, we overcome this problem by relying on an autonomous agent capable of combining the explanation obtained from the artist with the interpretation obtained by the machine learning mechanism. The artist provides the broader view, the historical context, and insights on the artist's background influencing his works (*e.g.,* political thought), while the XAI mechanism obtains the values of Wölfflin features and illustrates them on the transformed portrait of the user.

2.2 Styles in modern and contemporary art painting

The moment a painter takes out his brush and palette and starts painting on a canvas, it is possible that a new style will emerge [2]. Styles in painting had their heyday during the late 19th century and early 20th century. The style displayed or expressed in a painting can become part of an art movement, in which case there is a group of artists who have defined a certain style (in painting or other artistic disciplines). The individual interpretation was the most common impulse that made certain styles evolve to a peak, disseminate and transform into a new style. In the period mentioned above, also called 'modernism', we can count more than 100 different styles among which the Hurufiyya movement (Islamic calligraphy), Peredvizhniki (Russian 'wanderers' protesting against academic restrictions), Letras y figuras (depiction of letters of the alphabet during the Spanish colonial period in the Philippines) and many more [37].

Impressionism became one of the best-known movements in western Europe. The impressionists rejected, like their Russian contemporaries (the Peredvizhniki, Wanderers) classical and imperative aesthetic rules, incorporated by the Salons that detained a monopoly in the field of contemporary art exhibitions. An important innovator within this group was Claude Monet (1840 - 1926). In the American Magazine of Art (1927), Lilla Cabot Perry reveals the remarkable method Monet used to paint in the typical impressionistic style: He had grooved boxes filled with canvases placed at various points in the garden where there was barely room for him to sit as he recorded the fleeting changes of the light on his water-lilies and arched bridges. He often said that no painter could paint more than one half an hour on any outdoor effect and keep the picture true to nature, and remarked that in this respect he practiced what he preached [46].

Just after the turn of the century in 1907, Pablo Picasso painted *The Young Ladies of Avignon / The Brothel of Avignon*, where Aviñón refers to a street in Barcelona. This proto-cubist work introduces 'primitive style' from Africa in western art. At that time masks and sculptures from African countries became popular and were on sale in shops in Paris but also exhibited in museums. In the 'Autobiography of Alice B. Toklas', Toklas - who was befriended with Picasso - explains his radical turn in style as follows: "In these early days when he created cubism the effect of the African art was purely upon his vision and forms, his imagination remained purely Spanish" [52].

Ten years later. *Fountain*, an artwork of Marcel Duchamp (1887 - 1968) was submitted to the 1917 exhibition of the American Society of Independent Artists in New York. Sitting on a pedestal, turned upside down and signed with the artist's name R. Mutt. *Fountain* sparked raging controversy after Duchamp's colleagues refused to recognize the item as a legitimate work of art and requested it to be removed. Duchamp coined the term 'readymade' for this work but, while final evidence is missing, it could have been the 'Dada baroness' Elsa von Freytag - Loringhoven who came up with this idea and not Duchamp, who was a friend of her [19].

The utensil, called 'ready made' was an object, not made by an artist but ready to be chosen as an artwork. The case of the *Fountain* had the power to stop the development of all future styles and movements. This did not happen immediately but at the end of the Sixties. In *The conspiracy of Art*, the philosopher Jean Baudrillard describes the consequences of the removal of the artist from the artwork and even from the art world. He points to 'a new reality' that is not about creating styles but about creating interchangeable components that can each serve as 'reality': a procession of models providing autonomy for the virtual, freeing it from reality, and the simultaneous autonomy of reality that we now see functioning for itself - motu proprio - in a hallucinatory perspective, in other words, self - referential ad infinitum. Cast out from its own framework, from its own principle, extraneous, reality has itself become an extreme phenomenon. In other words, we can no longer think of it as reality, but only as otherworldly, as if seen from another world - as an illusion [5].

2.3 Neural Style Transfer

Style Transfer, often called Neural Style Transfer (NST), is the practice of manipulating a piece of data like an image or video to adopt the inherent style of another piece of data [29]. Common uses for this type of application come under the form of Deep Fakes, the image of a given person projected onto someone else's in a video, and the projection of a certain style onto a photograph, similar to a filter.

The main technology used nowadays for this type of functionality is the previously mentioned Generative Adversarial Network [21]. One of the more prominent GAN architectures with the objective of style transfer is the so-called Cycle-Consistent Generative Adversarial Network [61], CycleGAN in short, first proposed by Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. in 2017. One of the characteristic advantages of the CycleGAN is that it doesn't require paired training data to deliver results. What this means is that when a GAN is trained on two domains, by adding a piece of data to the set of one domain, it is not necessary to add the equivalent data of the opposite domain. This facilitates the creation of CycleGAN datasets since data can be compiled into each domain without further complications.

Just like with a regular GAN, both Generators and Discriminators are competing with each other. The Generators learn to create better fakes and the Discriminators learn to better detect the fakes. Together, the models get incrementally better at their tasks, learning from each other and adapting accordingly [8].

In addition to that, the Generator models are trained to not just create new images in the target domain, but instead reconstruct versions of the input images from the source domain. This is achieved by using generated images as input to the corresponding generator model and comparing the output image to the initial images. Passing an image through both generators is what we designate as a cycle. Together, each pair of generator models get trained to better reproduce the original source image, referred to as cycle consistency.

Another advanced implementation of style transfer can be found in StyleGAN [32]. In contrast to CycleGAN, this one does not need an entire dataset for the desired target-style application, one image proves sufficient. A StyleGAN works similarly to the traditional GAN implementation, however, after a StyleGAN has successfully been trained, it offers much finer control over the generated result than was previously possible. To achieve style transfer, StyleGAN is fed reference data on specific convolution layers. When it reaches those layers in the generation process, it will attempt to map the current sample to that target piece of data, effectively blending the image information into the end result.

2.4 XAI

The recent rapid development of Artificial Intelligence (AI) technology, as well as its widespread use in our daily lives, have raised several concerns about the human understandability of this sophisticated technology. To address this concern, eXplainable AI (XAI) [4] emerged to interpret the sometimes intriguing results of AI and ML learning mechanisms [22] as well as autonomous agents and robots [3]. In future AI systems, it is vital to guarantee a smooth human-agent interaction, as it is not straightforward for humans to understand the agent's state of mind, and explainability is an indispensable ingredient for such interaction [41].

Explanations can be provided by AI for a multitude of purposes including, control and debugging of AI systems, transparency, and accountability, as well as training and education. In the latter, case the main aim of the explanations is to help users understand how the system works and get a glance at its inner workings. Examples of these educational explanations including firefighter training [26], and UAV operation [43, 44]. Moreover, recently, the advantages of using XAI in humanities and arts. Yet, this line of research is in its early stages of development with many considerable challenges ahead [47].

3 The Smart Photobooth Project

The Smart Photobooth is a playful and interactive intelligent "machine" where the users can experiment with AI, and learn about the process of how intelligent machines are trained. For this, we have chosen to combine AI and Art. Science & Art have a vibrant and exciting relationship. Thus, they form a perfect combination to engage different audiences - of all ages, genders, and backgrounds - and entice them to learn more and explore computational methods. In particular, we have chosen to use AI to manipulate images to resemble artistic styles for two reasons: *(i)* humans are very visual and drawn to images, which are a powerful media to provoke interaction and enable easy communication — even more in the era of internet and social media - and *(ii)* the Smart Photobooth is similar to an interactive version of the Snapchat filter (which is widely known and popular). The Smart Photobooth is the perfect package to deliver our machine-learning & art message to a wider audience, in particular to teenagers, including those who are not technology drawn.

When the user enters the booth, they can take a portrait of themselves and then select a style to transform their image into. The set of available styles are selected by a professional artist (*e.g.,* impressionism, cubism, *etc.*). The style transfer is conducted using a NST mechanism (*cf.* Section 2). The NST mechanism is pre-trained using two training sets representing artistic styles as well as end-user portraits.

Once the user obtains the output (*i.e.,* their portrait pictured in the chosen style), they also obtain a multi-media presentation explaining both *(i)* the basic principles of the chosen style, its historical context, main contributors and most famous paintings, and *(ii)* an illustration of how the style influenced the visual features of the output image. This is obtained by the XAI mechanism which explains what of Wölfflin's features are present in the output image and how to they correspond to the style chosen by the user.

The project will be developed and presented in two different venues: a) Space 1 - workshops in the Scienteens Lab at the University of Luxembourg for STEM (science, technology, engineering, and mathematics) high school students in Luxembourg (April-July 2021); b) Space 2 - exhibition in the Luxembourg Science Center for various types of visitors - children, teenagers, families (August-December 2021). In 2022, we plan to have the Smart Photobooth exhibited for the whole year in the AI & Art Pavilion. The Pavilion, supported by Esch2022 Capital of Culture, will be providing various interactive programs and a series of exhibitions for all types of visitors for the duration of Esch2022.

4 Architecture

Figure 1 depicts our proposed architecture to explain the history of arts to human users by agents. On the bottom side of the figure, a training dataset is used to train the model after preprocessing it to handle any abnormalities. The result of the training model will be stored in the *Art Knowledge Base*, which will be the input to build the deployment machine learning model. The GAN-based NST algorithm takes as input: *(i)* The input image of the human user to be converted; *(ii)* an input dataset of images similar to the human user image; *(iii)* the deployment model features. The output of the algorithm is the output image with a specific art style. This image will be presented to the human and will help along with some features from the deployment model to provide some ML analysis to be used for the explanations. The explanation model in the agent formulates the explanations based on three sources:

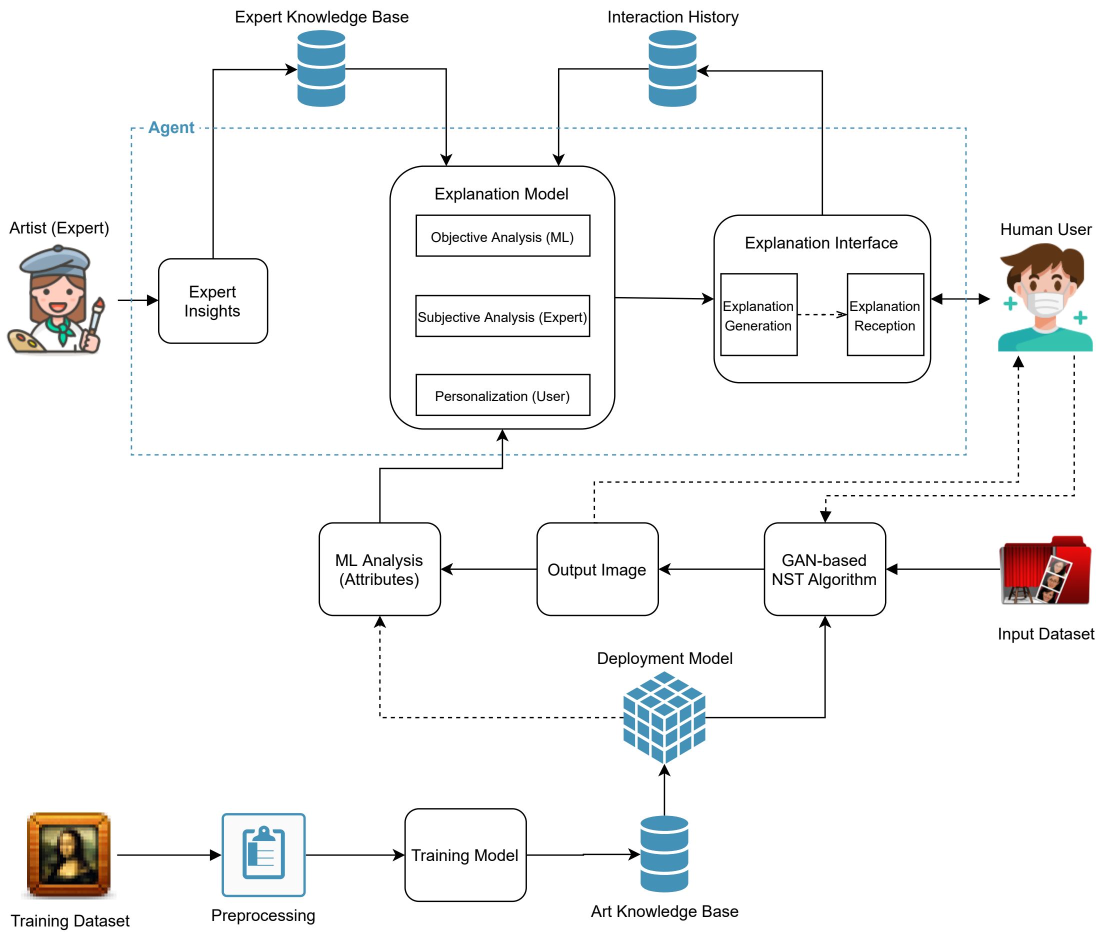

Fig. 1. Explainable Human-Agent Architecture for Art: Bottom side is the machine learning model including the *GAN-based NST Algorithm*. Top side is the agent model responsible for explainability. Two-ways arrow between the *Explanation Interface* and *Human User* to highlight an interaction where the feedback from the user is used to update the *Explanation Model*.

(i) An objective analysis based on the insights provided by the ML algorithm and model; *(ii)* A subjective analysis provided by the expert in the domain (the artist) which will store relevant knowledge in the *Expert Knowledge Base*; *(iii)* Another subjective analysis performed on the human user side to guarantee that the personalized explanations are built based on the preference of the human. For this step, the literature highlights the need to move towards human–centered Models as explanations are subjective [51]. This knowledge about the human user is collected depending on the various interactions with him/her and stored in the *Interaction History Base*, hence the two-ways arrow between the *Explanation Interface* and the *Human User*.

Finally, the formulated explanations are communicated by the agent to the human user through an interface that allows for the following two tasks: *(i)* generating the explanations based on the explanation components formulated in the explanation model. The generation could use templates or be more sophisticated by relying on NLP techniques. *(ii)* Providing the explanation to the human through the explanation reception process that allows the interaction with the human and considers its cognitive load.

5 Open Challenges and Research Directions

Using XAI for AI & Art dissemination is a domain in its early stages of development with most of the pioneering work carried out at the conceptual front ([15]). The Smart Photobooth is a work-in-progress aiming to operationalize and test XAI-based art dissemination in a real-life context. As explained in the previous sections, the project proposes an architecture that combines expertise from the artist, input from the end-user, and the output of the most advanced machine learning mechanisms. The intelligent agent proposed by our architecture is in charge of combining these heterogeneous data and making them accessible and understandable by different stakeholders. The project still being in its early implementation phase, this section identifies the following challenges and research directions we will pursue to address them.

(i) **Explaining heterogeneous AI systems**: The Smart Photobooth project involves multiple AI systems. Namely: The NST systems, the style interpretation system, and the agent which is in charge of obtaining the artist's explanation, and the input characterizing the user. Combining these heterogeneous explanations is a challenging task notably because it involves both symbolic knowledge (user data and artist input) and sub-symbolic knowledge within the black-box machine learning mechanism. One potential solution to address this issue is to resort to the latest advances in neuro-symbolic AI which proposes to integrate the symbolic AI systems [20] (*i.e.,* agents knowledge is represented by logic and reasoning) with the sub-symbolic knowledge within the GANs and the convolutional neural networks. Recent works in XAI suggested that this neuro-symbolic integration is highly beneficial to XAI [9]. In particular, compared with the current approach which relies on a simple concatenation of expert explanations with the visual descriptors originating from the neural network, the neuro-sym-

bolic approach allows for a better understanding of the system's output since the symbolic knowledge extracted from the neural network can be reasoned and manipulated by the agent. Moreover, information extracted also be made accessible/understandable by the artist who can also interact with the knowledge to tune the performance of the machine learning mechanism. *(ii)* **Users with Special Needs:** Most of the explanations and insight provided by the Smart Photobooth are communicated audio-visually. For this reason, the content is inaccessible for users with visual and/or audio impairment. One alternative solution is to rely on solutions currently being developed to make cultural heritage accessible for this category of users. For instance, *bind photography* [59] can assist users with low visual acuity take their photos by providing audio feedback that facilitates aiming the camera. Additionally, online accessibility of explanations is needed for those who could not visit the artworks galleries [42] *(iii)* **Gamification:** In recent years, gamification [14, 25, 28] has been a trending topic in domains as diverse as education, information studies, and human–computer interaction [50]. It is also considered an interesting means of supporting user engagement and enhancing positive use patterns, such as increasing user activity, social interaction, or quality and productivity of actions [23]. These desired use patterns are considered to emerge as a result of the positive, intrinsically motivating [49], and gameful experiences [28] brought about by game/motivational affordances implemented into a service [24]. Recently, art organizations have also sought to gamify different aspects of their institutions to engage visitors, increase fundraising, or improve marketing objectives [6]. The application of gamification in education and outreach settings like in the context of this paper is still a relatively new trend, but it has gained attention due to its ability to increase student motivation and engagement. According to [16], there are three main concerns when considering gamifying the learning experience: "*(i)* insufficient evidence exists to support the long-term benefits of gamification in educational contexts; *(ii)* the practice of gamifying learning has outpaced researchers' understanding of its mechanisms and methods; *(iii)* the knowledge of how to gamify an activity in accordance with the specifics of the educational context is still limited."

6 Conclusions

This paper presented the Smart Photobooth, an interdisciplinary outreach and knowledge dissemination project organized by the University of Luxembourg. The project relies on XAI to explain the history of art and artistic styles to end-users. Delivering such explaining requires a combination of explanations provided by the human expert (an artist) with interpretation obtained from a machine learning mechanism. To achieve this synergy, the paper proposed an architecture powered by an agent who is in charge of accomplishing this combination. The components of the architecture were identified and explained. Open issues and challenges were identified with their potential solutions. The Smart Photobooth project is a work-in-progress, as the project started in early 2021. Currently, the Photobooth is being installed and the proposed architecture and the XAI mechanism are being implemented. The next step is to evaluate the explanations and assess how well they performed in engaging the end-users and enhancing their knowledge on the history of art. To do so, specific XAI metrics [27] will be defined, several user studies will be conducted, and their results will be statistically validated, studied, and analyzed.

Acknowledgment

The Smart Photobooth project is funded by FNR (Fond National de Recherche), Luxembourg. Reference: PSP-C2020-2/15417971.

References

1. Images of artworks in the public domain. https://www.metmuseum.org/about-the-met/policies-and-documents/image-resources.
2. Svetlana Alpers. Style is what you make it: the visual arts again. *The Concept of Style, ed. Berel Lang (revised and expanded edition, Ithaca and London: Cornell Univ. Pr., 1979, 1987)*, pages 137–62.
3. Sule Anjomshoae, Amro Najjar, Davide Calvaresi, and Kary Främling. Explainable agents and robots: Results from a systematic literature review. In *18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, pages 1078–1088. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
4. Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, et al. Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information Fusion*, 58:82–115, 2020.
5. Jean Baudrillard. *The conspiracy of Art.* Manifestos, Interviews, Essays, Semiotext(e) Foreign Agents Series, 2003.
6. Andre Bouchard. Gamification in the arts. 2013.
7. Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2018.
8. Jason Brownlee. How to develop a cyclegan for image-to-image translation with keras, Aug 2019.
9. Roberta Calegari, Giovanni Ciatto, and Andrea Omicini. On the integration of symbolic and sub-symbolic techniques for xai: A survey. *Intelligenza Artificiale*, 14(1):7–32, 2020.
10. Marcello Carrozzino and Massimo Bergamasco. Beyond virtual museums: Experiencing immersive virtual reality in real museums. *Journal of Cultural Heritage*, 11(4):452–458, 2010.
11. Eva Cetinic, Tomislav Lipic, and Sonja Grgic. Learning the principles of art history with convolutional neural networks. *Pattern Recognition Letters*, 129:56–62, 2020.
12. Eva Cetinic and James She. Understanding and creating art with ai: Review and outlook. *arXiv preprint arXiv:2102.09109*, 2021.
13. António Coelho, Pedro Cardoso, Maria van Zeller, Liliana Santos, José Raimundo, and Roberto Vaz. Gamifying the museological experience. 2020.

14. Sebastian Deterding, Dan Dixon, Rilla Khaled, and Lennart Nacke. From game design elements to gamefulness: defining" gamification". In *Proceedings of the 15th international academic MindTrek conference: Envisioning future media environments*, pages 9–15, 2011.
15. Natalia Díaz-Rodríguez and Galena Pisoni. Accessible cultural heritage through explainable artificial intelligence. In *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, pages 317–324, 2020.
16. Darina Dicheva, Christo Dichev, Gennady Agre, and Galia Angelova. Gamification in education: A systematic mapping study. *Journal of Educational Technology & Society*, 18(3):75–88, 2015.
17. Ahmed Elgammal, Bingchen Liu, Diana Kim, Mohamed Elhoseiny, and Marian Mazzone. The shape of art history in the eyes of the machine. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
18. Emma Frid, Hans Lindetorp, Kjetil Falkenberg Hansen, Ludvig Elblaus, and Roberto Bresin. Sound forest: Evaluation of an accessible multisensory music installation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2019.
19. Irene Gammel. *Baroness Elsa, Gender, Dada and Everyday Modernity. A Cultural Biography*. MIT press, 2003.
20. Artur S d'Avila Garcez, Krysia B Broda, and Dov M Gabbay. *Neural-symbolic learning systems: foundations and applications*. Springer Science & Business Media, 2012.
21. Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. Jun 2014.
22. Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5):1–42, 2018.
23. Juho Hamari. Transforming homo economicus into homo ludens: A field experiment on gamification in a utilitarian peer-to-peer trading service. *Electronic commerce research and applications*, 12(4):236–245, 2013.
24. Juho Hamari, Jonna Koivisto, and Harri Sarsa. Does gamification work?–a literature review of empirical studies on gamification. In *2014 47th Hawaii international conference on system sciences*, pages 3025–3034. Ieee, 2014.
25. Juho Hamari and Vili Lehdonvirta. Game design as marketing: How game mechanics create demand for virtual goods. *International Journal of Business Science & Applied Management*, 5(1):14–29, 2010.
26. Maaike Harbers, Karel Van Den Bosch, and John-Jules Meyer. A methodology for developing self-explaining agents for virtual training. In *International Workshop on Languages, Methodologies and Development Tools for Multi-Agent Systems*, pages 168–182. Springer, 2009.
27. Robert R Hoffman, Shane T Mueller, Gary Klein, and Jordan Litman. Metrics for explainable ai: Challenges and prospects. *arXiv preprint arXiv:1812.04608*, 2018.
28. Kai Huotari and Juho Hamari. Defining gamification: a service marketing perspective. In *Proceeding of the 16th international academic MindTrek conference*, pages 17–22, 2012.
29. Yongcheng Jing, Yezhou Yang, Zunlei Feng, Jingwen Ye, Yizhou Yu, and Mingli Song. Neural style transfer: A review. Oct 2018. arXiv: 1705.04058.
30. Yongcheng Jing, Yezhou Yang, Zunlei Feng, Jingwen Ye, Yizhou Yu, and Mingli Song. Neural style transfer: A review. *IEEE transactions on visualization and computer graphics*, 26(11):3365–3385, 2019.

31. Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019.
32. Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. Mar 2019.
33. Galena Kostoska, Marcos Baez, Florian Daniel, and Fabio Casati. Virtual, remote participation in museum visits by older adults: a feasibility study. In *8th International Workshop on Personalized Access to Cultural Heritage (PATCH 2015), ACM IUI 2015*, pages 1–4, 2015.
34. Galena Kostoska, Arnold POS Vermeeren, Joke Kort, and Charlie Gullström. Video-mediated participation in virtual museum tours for older adults. In *10th International Conference on Design & Emotion, 27-30 September 2016, Amsterdam*. The Design & Emotion Society,, 2016.
35. Rajat Kulkarni, Rutik Gaikwad, Rudraksh Sugandhi, Pranjali Kulkarni, and Shivraj Kone. Survey on deep learning in music using gan. *International Journal of Engineering Research & Technology*, 8(9):646–648, 2019.
36. Alberto Lamas, Siham Tabik, Policarpo Cruz, Rosana Montes, Álvaro Martínez-Sevilla, Teresa Cruz, and Francisco Herrera. Monumai: Dataset, deep learning pipeline and citizen science based app for monumental heritage taxonomy and classification. *Neurocomputing*, 420:266–280, 2021.
37. Berel Lang. *The concept of style*. Cornell University Press, 1987.
38. Germain Latour. *Guernica, histoire secrète d'un tableau*. Média Diffusion, 2013.
39. Negrevergne Benjamin Lecoutre, Adrian and Florian Yger. Recognizing art style automatically with deep learning. In *Proceedings of Machine Learning Research*, number 77, pages 327–342, 2017.
40. Hui Mao, Ming Cheung, and James She. Deepart: Learning joint representations of visual arts. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 1183–1191, 2017.
41. Yazan Mualla. *Explaining the Behavior of Remote Robots to Humans: An Agent-based Approach*. PhD thesis, University of Burgundy - Franche-Comté, Belfort, France, 2020. 2020UBFCA023.
42. Yazan Mualla, Timotheus Kampik, Igor H. Tchappi, Amro Najjar, Stéphane Galland, and Christophe Nicolle. Explainable agents as static web pages: Uav simulation example. In *Explainable, Transparent Autonomous Agents and Multi-Agent Systems*, pages 149–154, Cham, 2020. Springer International Publishing.
43. Yazan Mualla, Amro Najjar, Timotheus Kampik, Igor Tchappi, Stéphane Galland, and Christophe Nicolle. Towards explainability for a civilian uav fleet management using an agent-based approach. *arXiv preprint arXiv:1909.10090*, 2019.
44. Yazan Mualla., Igor Tchappi., Amro Najjar., Timotheus Kampik., Stéphane Galland., and Christophe Nicolle. Human-agent explainability: An experimental case study on the filtering of explanations. In *Proceedings of the 12th International Conference on Agents and Artificial Intelligence - Volume 1: HAMT,*, pages 378–385. INSTICC, SciTePress, 2020.
45. Xin Ning, Fangzhe Nan, Shaohui Xu, Lina Yu, and Liping Zhang. Multi-view frontal face image generation: A survey. *Concurrency and Computation: Practice and Experience*, page e6147, 2020.
46. LillaCabot Perry. Reminiscences of claude monet from 1889 to 1909. *The American Magazine of Art*, XVIII(3):123, March 1927.
47. Galena Pisoni, Natalia Díaz-Rodríguez, Hannie Gijlers, and Linda Tonolli. Human-centred artificial intelligence for designing accessible cultural heritage. *Applied Sciences*, 11(2):870, 2021.

48. Pauliina Raento and Cameron J Watson. Gernika, guernica, guernica?: Contested meanings of a basque place. *Political Geography*, 19(6):707–736, 2000.
49. Richard M Ryan and Edward L Deci. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American psychologist*, 55(1):68, 2000.
50. Katie Seaborn and Deborah I Fels. Gamification in theory and action: A survey. *International Journal of human-computer studies*, 74:14–31, 2015.
51. Ronal Singh, Paul Dourish, Piers Howe, Tim Miller, Liz Sonenberg, Eduardo Velloso, and Frank Vetere. Directive explanations for actionable explainability in machine learning applications. *arXiv preprint arXiv:2102.02671*, 2021.
52. Gertrude Stein. *Autobiography of Alice Toklas 1907 - 1914*. The Library of America, 1933.
53. Gjorgji Strezoski and Marcel Worring. Omniart: a large-scale artistic benchmark. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 14(4):1–21, 2018.
54. Wei Ren Tan, Chee Seng Chan, Hernán E Aguirre, and Kiyoshi Tanaka. Ceci n'est pas une pipe: A deep convolutional network for fine-art paintings classification. In *2016 IEEE international conference on image processing (ICIP)*, pages 3703–3707. IEEE, 2016.
55. Mukhiddin Toshpulatov, Wookey Lee, and Suan Lee. Generative adversarial networks and their application to 3d face generation: A survey. *Image and Vision Computing*, page 104119, 2021.
56. Roberto Trotta, Daniel Hajas, José Eliel Camargo-Molina, Robert Cobden, Emanuela Maggioni, and Marianna Obrist. Communicating cosmology with multisensory metaphorical experiences. *Journal of Science Communication*, 19(2), 2020.
57. Heinrich Wölfflin. *Principles of art history*. na, 1915.
58. Suiyun Zhang, Zhizhong Han, Yu-Kun Lai, Matthias Zwicker, and Hui Zhang. Stylistic scene enhancement gan: mixed stylistic enhancement generation for 3d indoor scenes. *The Visual Computer*, 35(6):1157–1169, 2019.
59. Yuhang Zhao, Shaomei Wu, Lindsay Reynolds, and Shiri Azenkot. The effect of computer-generated descriptions on photo-sharing experiences of people with visual impairments. *Proceedings of the ACM on Human-Computer Interaction*, 1(CSCW):1–22, 2017.
60. Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.
61. Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. Aug 2020.