Crea.visions: A Platform for Casual Co-Creation with a Purpose Envisioning the Future through Human-AI Collaboration with Multiple Stakeholders

Janet Rafner[1], Blanka Zana[1], Tristan Beolet[1], Safinaz Büyükgüzel[1], Ewen Michel[1], Neil Maiden[2], Sebastian Risi[3], Jacob Sherson[1]

[1] Center for Hybrid Intelligence, Aarhus University, Aarhus, Denmark
[2] Digital Creativity at the Bayes Business School, University of London, London, England
[3] Robotics, Evolution and Art Lab (REAL), IT University of Copenhagen, Copenhagen, Denmark

janetrafner@mgmt.au.dk, blanka.zana@gmail.com, tristan.beolet@student-cs.fr, safinaz@mgmt.au.dk, ewen.michel@student-cs.fr, Neil.Maiden.1@city.ac.uk, seb@itu.dk, sherson@mgmt.au.dk

Abstract

With recent advances in Artificial Intelligence and increasing emphasis on human augmentation and collaboration, time is ripe for AI-enhanced support tools which empower the public to formulate and visualize a collective vision of societal issues such as climate change. Here, we report on crea.visions, a platform for human-AI co-creation within Sustainable Development Goals centered community engagement. We present in-the-wild experiments with four versions of crea.visions involving 1,000+ participants and 25,000+ generated images over three years: Versions 1 and 2 focused on developing the novel tool empowering citizens to artistically communicate their favorite abstract societal issues. In versions 3 and 4, the generic image generation GAN was replaced by custom-trained versions for Venice and Paris respectively. Refining the platform towards community-specific action, users of version 4 can geotag their identified problems, submit solutions candidates, and are actively linked up with relevant NGOs. Finally, version 4 includes the first workflow todate which combines AI image-generating modalities of sliders and text-to-image.

Introduction

The OECD (Schleicher, 2010), the World Economic Forum (Belsky, 2020) and many other international agencies argue that creativity is one of the top five skills in the 21st century. This includes how creativity is necessary for formulating and implementing a wide range of local and global solutions to complex issues such as the Sustainable Development Goals.[1] Thus, methods for fostering, improving, and facilitating human creativity have been studied for decades ranging from ideation interventions Baas, De Dreu, and Nijstad (2008); Santanen et al. (2004) to creativity support tools (Frich et al., 2019; Sielis, Tzanavari, and Papadopoulos, 2009). However, with the ever-improving Artificial Intelligence (AI) technologies, how AI can augment human creativity is becoming a prominent area for research and development. Human-AI co-creative systems, involve at least one human agent and one artificially intelligent agent collaborating with each other to build creative artifacts (Davis, 2013; Kantosalo and Toivonen, 2016). This collaborative activity has been defined as mutually influential contributions (Davis, 2013), mixing human and computational initiatives (Yannakakis, Liapis, and Alexopoulos, 2014) and the sharing of creative responsibility (Kantosalo and Toivonen, 2016).

A recent trend called Casual Creators can be considered a subcategory of human-AI co-creative systems. Casual Creators promote efficient and enjoyable exploration of a possibility space, leading to the creation of unexpected artifacts that inspire feelings of pride, ownership, and creativity in the users who create them (Compton and Mateas, 2015). At the core of these products often lie AI generators which empower amateur creativity by mapping simple low-dimensional input domains, e.g., sliders, to a complex high-dimensional output domain, e.g., images (Gajdacz et al., 2021). Beyond traditional entertainment purposes, researchers have demonstrated the utility of casual creators for creativity assessment (Rafner et al., 2020; Rafner, 2021; Gajdacz et al., 2021), raising awareness of societal issues (Chang and Ackerman, 2020; Luccioni et al., 2021; Rafner et al., 2021) and visualizing the future (Epstein, Schroeder, and Newman, 2022). These tools predominantly use Generative Adversarial Networks (GANs), a well-known machine-learning model designed to produce artificial images that are nearly indistinguishable from real images (Borji, 2019). GAN's abilities present two major advantages when it comes to designing co-creative systems: the generation is deterministic, i.e., non-random, and the generative part of the network is continuous and locally coherent over the latent space, [2] both properties making it in-

[1] https://sdgs.un.org/goals

[2] n-dimensional vector space from which the noise vectors used for image generation are sampled. Notions of latent space and de-

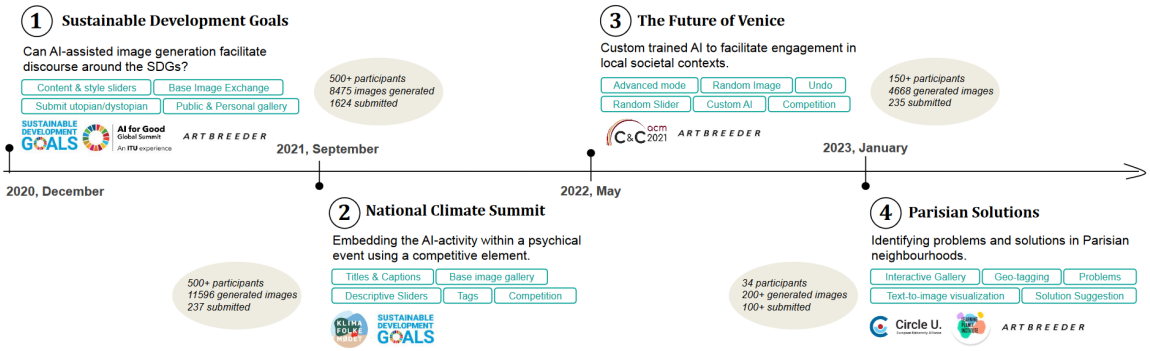

Figure 1: This figure presents a timeline of versions of crea.visions, the context in which it was used, the new key elements, stake holders, as well as number of participants and images generated and submitted for each of the four events.

tuitive and controllable to explore for naive users.

Despite these promising developments, there remains a significant research gap. Current literature on casual creators used for research purposes is still in its infancy; the articles cited above primarily focuses on prototype descriptions or pilot studies of user interactions, with little attention given to large-scale multi-stakeholder studies (e.g., participants, government officials, companies) or successive iterations of a tool. In other words, the potential of Casual Creators to contribute to large-scale, public use for various research and societal applications remains largely unexplored and untested in real-world settings.

To address this gap, we propose an approach grounded in the Community Citizen Science (CCS) framework (Hsu and Nourbakhsh, 2020). Our objective is to move beyond prototype and pilot studies, and extend the application of Casual Creators into large-scale, ecological environments with diverse stakeholders. The CCS framework, emphasizing participatory democracy and community co-design, offers a promising route towards achieving this goal. In the context of Human-Computer Interaction, CCS has been utilized to enhance scientific research and community empowerment (Hsu and Nourbakhsh, 2020). Examples include mobile apps that allow residents to track pollution odors or generate high-resolution landscape imagery (Hsu and Nourbakhsh, 2020). However, the application of CCS principles to human-AI co-creative systems remains largely unexplored. This study aims to address this uncharted territory, presenting the first known application of CCS principles in this context. We believe this approach can yield significant benefits, opening up new possibilities for the use and advancement of Casual Creators.

This article presents the development of four versions of crea.visions, a platform for human-AI co-creation designed for community involvement and social good; see Fig. 1 for an overview on the four versions. We aim to show how combining human-AI co-creative technology and CCS principles holds potential for new forms of public engagement on complex socio-environmental topics such as climate change. We approached the development of crea.visions with design goals (DG) of a user-friendly and intuitive interface (DG1), meaningful content and motivation for activities (DG2), meeting technical requirements (DG3), and community buy-in and co-design with multiple stakeholders (DG4). We provide an overview of the platform, describe and analyze data from four experiments over three years. We conclude with a discussion on future directions for crea.visions.

coder/generator tend to be used in similar contexts.

Crea.visions System Overview

The core mechanics of crea.visions allows users to blend 'style' (small scale features and texture) and 'content' (large scale features) components of a set of source images into new images (see Fig. 2). Images are generated by the pre-trained StyleGAN2 (Karras et al., 2020) and the system is developed in the unity game engine.

Crea.visions enables a form of alternating, turn based, co-creativity where first the human provides an input and then the model generates an image based on the users input. The interface is simple with minimal features to allow for very brief on-boarding instructions before users can generate contextually meaningful images.

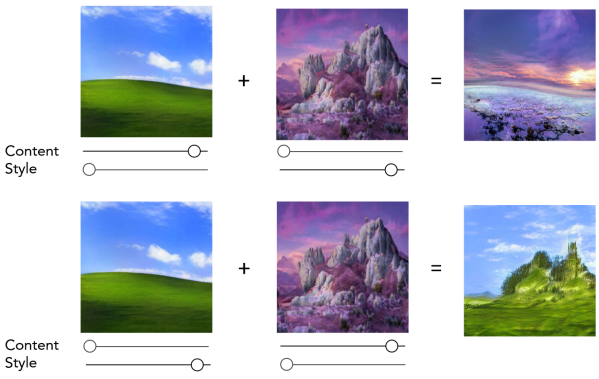

Figure 2: This image illustrates how in crea.visions, players use two sliders per image in order to control the 'style' (small scale features and texture) and 'content'(large scale features) they want each image to contribute into a new image.

Crea.visions Version 1: Sustainable Development Goals, Online Data Collection

The main purpose of V1 was to demonstrate the technical viability of crea.visions (DG3) and the problem framing needed to engage the public in creating thought-provoking visions of utopias and dystopias in order to raise awareness of socioscientific problems related to the SDGs. In this experiment, as with the others, IRB approval was received from Aarhus University. Participants gave informed consent prior to participation.

Multi-Stakeholder Alignment and Experimental Design

V1 was developed with Artbreeder, a popular online image generation platform, and the United Nations online platform, AI4good. The authors met regularly over the course of six months with the representatives from both platforms to optimize usability. Due to Covid19 the launch was online only. The use of crea.visions was promoted by all stakeholders via social media channels including Twitter, Facebook, LinkedIn, and Instagram. Participation was completely voluntary. Data was collected between December 2020 and March 2021 and included slider movements and user clicks, and submitted images.

Technical Considerations

The preexisting StyleGAN model (Karras et al., 2020) was used in V1. The roughly 200 base images originally from the Artbreeder community were selected by the first author with the criteria of diversity in style, content, color, and motif to facilitate a wide range of possible utopian and dystopian images. The V1 experiment also served to test the capacity of the tool (i.e., number of simultaneous users) as image blending requires heavy computational power.

Gameplay and User Experience

Participants were presented with five options on the landing page: a tutorial, image blending, a public gallery, next steps (project information), and a personal gallery. Participants were given the challenge to imagine the world 50 years from now and generate images of possible futures. Given the content of the images, they were particularly effective for climate related issues (fires, flooding, etc.). The user flow guided first-time users to begin with the tutorial, explore the public gallery for inspiration, then proceed to blend up to four images which they could change by clicking on them. A click on a base image randomly loaded a new one. When a player finished creating their image, they could submit the blended image, tagging it as either utopian or dystopian. The submitted images were anonymously published in a public gallery for voting (Fig 3, Right) and could be downloaded or further edited by other participants. The personal gallery stored all images created during the individual session.

Descriptive Statistics, Exploratory Analysis, and Lessons Learned

In total there were 580 user sessions and 8475 images generated. Exploratory analysis from this study is published in Rafner et al. (2021), showing that V1 allowed users to create images that express both anxiety and hope for the future, affirmed that user-generated images express these ideas in ways that are meaningful to others, and began to investigate which specific features of images (color, motif, style) are more closely related to dystopian or utopian ideas. The research explored in Rafner et al. (2021) presents results on image analysis of the generated images from V1 where as the present article focuses much more on the system design user experience, and multi-stake holder alignment for V1-V4. The testing of V1 also identified that one virtual machine with an Nvidia A100 GPU, 4 cores and 25GB RAM could support approximately 50 simultaneously users. The partners were happy with the outcome of the launch as a proof of concept, however as suggested by the CCS Framework, future iterations needed to focus on co-creating features with multiple stakeholders (DG4), contextualizing the activities to bring more meaning and adding more options for the participants for expressing intentions (DG2).

Crea.visions Version 2: National Climate Summit: In-Person Events with Students and the General Public

The main purpose of V2 was to improve user control, specifically participants' ability to describe their intentions (DG1) and to evaluate if a competition component was beneficial for crea.visions player motivation (DG2).

Multi-stakeholder Alignment and Experimental Design

In September 2021, V2 was launched at a Danish People's Climate Summit[3] in Middelfart, Denmark; two data collection sessions were conducted. Participants were instructed to create their vision of the future 50 years from now using V2. Log data, titles, captions, and tags were collected and participants provided qualitative feedback. Participants were incentivized to relate their work to environmental topics, in particularly climate change, while also specifically "having to picture their view of the future." The first data collection was with 500 high school students who were participated during in-class time, working in teams to produce over 170 submissions over one week; students used their own laptops. The authors met with coordinators from the school multiple times prior to the event to prepare an introductory presentation on human-AI co-creativity for the students. The second data collection took place at the public climate summit event, which hosted numerous environment-related speakers and activities. Participants interacted with V2 on laptops, available at open stands (Fig 4). Participation in the event was advertised via social media by all partners and participation was on a voluntary basis. However, participants could submit their creations to a competition with a prize of 1350 EUR. Submissions were evaluated by an expert panel based on the novelty and appropriateness. Titles and captions played a crucial role in the competition, as they

[3] Klimafolkemødet at https://klimafolkemoedet.dk/

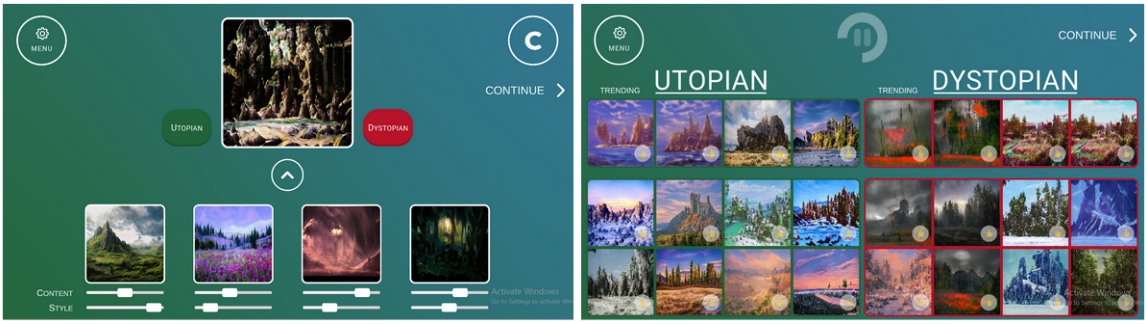

Figure 3: V1. Left: Blending interface. Right: Public gallery of utopian and dystopian images.

Figure 4: V2: Photos from climate summit event.

allowed the intentions of the player to be described (Fig 5, Right).

At the climate summit's closing ceremony, the mayor of Middlefart presented the award. The winning submission was made by a group of three students, who commented "It was fun to try something new where you could express yourself creatively." A highschool teacher part of the organizing team said that crea.visions was "very user-friendly, and quite easy to navigate" and that "for me, and for my colleagues, I think the most positive thing has been that it is so different from what we are used to in everyday life". The leader of the climate summit concluded that "we had a really good collaboration with the team of crea.visions, and I expect that we will continue with next year's meeting".

Technical Considerations

The GAN used and image selection was identical to V1.

Gameplay and User Experience

Participants were first shown a video which explained the game and competition. They then proceed to blend images. V2 was equipped with a two-step submission page that allowed participants to provide a title, a description and tags (Fig 5, Left), and rate their image properties on various scales, (e.g. natural/man-made, permanent/reversible - see Fig. 6). The new submission process was designed to prompt deeper reflection in the players regarding various dimensions of the image, and add narrative to the submitted images. The base image swap feature was developed into a controlled selection process, where upon a click on a base image a gallery of images would open, allowing people to browse through various options (Fig. 6) as opposed to the random cycling through of images, supported by V1. Similarly to V1, V2 had a social library function and personal user gallery.

Descriptive Statistics, Exploratory Analysis, and Lessons Learned

In total, 11 596 images were generated and 237 were submitted to the competition. Exploratory data analysis was performed to investigate how participants explored the possibility space before submitting an image (e.g., are they thoughtful creations or simply random submissions) to help determine if V2 was appropriately scaffolding the interaction (DG1). Users generated on average 43 images before submitting, with a range of 0 to 272. The high number of generated images before submission was interpreted as a proxy for engagement: if the submission was preceded by a high number of iterations, it could be considered a thoughtful creation. On average, the highest ranked image creators (N=16) explored more than other players, however as the sample size of the winners group was small, no significant conclusions can be drawn regarding group differences. Qualitative feedback indicated it was difficult for the players to know how much of the image space they had explored and were frustrated that they were unable to go back to previously generated images. By performing sentiment analysis on titles and captions, we found that people's dystopian/utopian slider usage was positively correlated ($r=.24$, $p=.009$) with the sentiment of their captions (e.g., the more utopian an image was rated on the scale, the more positive the caption's sentiment was), indicating player intentions are reflected through both slider use and captions. It was observed that a thoughtful submission took at least 20mn. The gallery feature and the 'upvoting' was intended to provide a shorter, yet meaningful interaction for those with less time, but qualitative feedback indicated this was not the case (DG3).

Logistically speaking we also observed that in order to make a thought-through submission, participants took at

Figure 5: V2. Left: image blending interface Middle: The winner receiving the prize. Right: The winning submission.

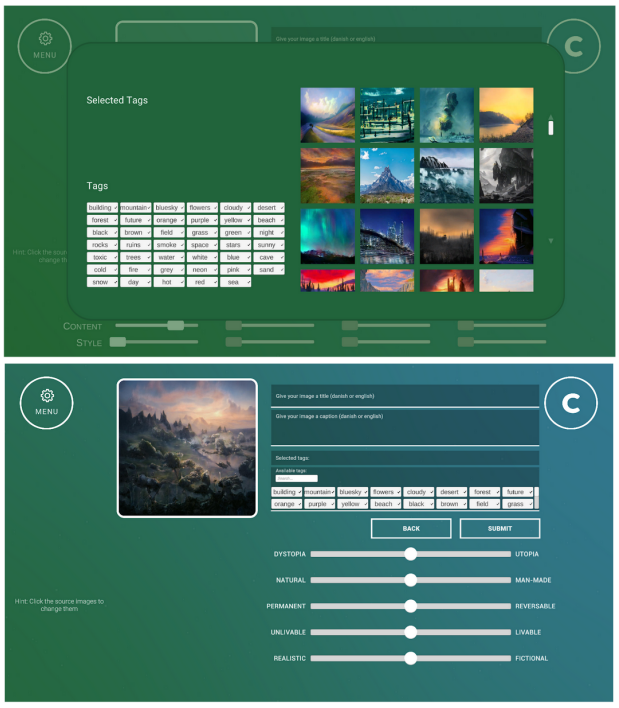

Figure 6: Screenshots from V2. Top: Base image Gallery. Bottom: Second Submission Page.

least twenty minutes and required a comfortable place to sit and concentrate, out of the sun to avoid glare on their screens.

Crea.visions Version 3: The Future of Venice: Venetian Residents and Tourists

The main purpose of V3 was to use custom images for local specificity (DG2, DG4) and usability improvements to allow the participants to navigate through the GAN space more fluently (DG1, DG3).

Multi-Stakeholder Alignment and Experimental Design

V3 was developed for the 2022 Creativity & Cognition Conference in Venice, Italy. As part of the crea.visions week event, pop-up workshops were held on the streets of Venice. Log data from the blending, the titles, and captions was saved. Participants were asked to provide general qualitative feedback about their experience. Participants' ages ranged from 16-70+. Through social media, word of mouth, and on-the-street recruitment, the public was invited to create images of the future of the city in order to spark dialogues and self-reflection around sustainability, climate change, and tourism. As recommended by the CCS framework we coordinated with our local collaborators, including the General Co-Chair of the conference; we visited locations that attract a variety of groups, such as cafes, schools, parks, and local communities. Participation was voluntary, but participants were provided with refreshments while they sat and generated the images. Participants could also choose to submit their visions (images with descriptions) to a competition; winning visions were exhibited from 22-25 June 2023 as part of the conference art exhibition. Visions were evaluated in the same manner as V2. At the exhibition, we provided printed-out take-away postcards of 16 selected images.

Figure 7: V3: Pictures from the Venice event.

Technical Considerations

Unlike V1 and V2 which used a generic pre-trained GAN, our Venetian partners wanted a version that produced recognizable Venetian architectural features to be contextually meaningful to the target group (DG2, DG4). Thus, the StyleGAN model was retrained (Varkarakis, Bazrafkan, and Corcoran, 2020) on images of Artbreeder and Venetian buildings (see Fig. 9). In order to maintain the diversity of the pre-trained StyleGAN model, and achieve the continuous building-like latent space needed for this application, the standard training procedure for StyleGAN was used, even though it is computing intensive and required high-end hardware. The model requires 12 GB of VRAM, and common consumer-grade graphic cards often have between 4 and 8 GB of VRAM so a fifth of a Nvidia A100 has been used

Figure 8: Left:V3 blending interface, advanced mode toggled on. Middle: Photograph from crea.visions Week. Right: selected images in postcard format including the final image, source images, title, caption and creators name.

for this purpose. The retraining set was composed of 5026 images, 1855 of which were crowdsourced photos taken of Venice by our local partners. The rest of the images were generated by the original model. To bias the model to produce more images with certain features (vegetation, water, etc.) 16 base images containing the desired features were selected by the first author. These were blended together in order to achieve a reasonably sized dataset with sufficient variations:

Figure 9: Image illustrating the GAN's blending process for training image generation in V3. The far left and right images are the source images, and the images in the middle represent images that could be blended using these two images with various slider settings.

Gameplay and User Experience

Participants received a verbal introduction to the tool and then blended images using laptops pre-loaded with V3. In V3, the blended image is in the middle of the interface, and the source images are in the corners (Fig. 8, Left) to accommodate the advanced mode (described below). After blending images, participants moved to the submission page, giving their image a title, caption, and tags. The base image selection gallery was carried on from V2. New features were added, such as a random images button, which set four random base images and a random sliders button changed all four images' slider settings in order to facilitate the participants' exploration process (DG1). An undo button was added, allowing the user to return to the previous image settings as well as a button enabling the advanced mode (See Fig. 8, Left), displaying 16 images around the blended image, each the result of slightly adjusted slider settings compared to the current image. These served as a preview of possible images. Once a preview image was clicked on, the respective sliders would change and the selected image would be generated.

Descriptive Statistics, Exploratory Analysis, and Lessons Learned

4668 images were generated and 235 submitted. We were pleasantly surprised by the positive reception of V3 by Venetians aged 60+, many of whom voiced that they used little technology. A Venetian marketing student said, "It can give the people a great impression of the future. Not just an image in your mind, but something you can see on the laptop. So you can create it and actually see it." Multiple participants mentioned that the platform was "easy and intuitive" however, image content variety (more people, vegetation) and aspects of usability (text size, tool compatibility) could be improved. During the sessions, it became apparent that the two-step submission process did not lend itself intuitively to the generation of multiple submissions. The introduction of the second screen made the submission process feel rather final and gave people the impression that they finished a task. Additionally, due to the additional computational load of generating the 16 supporting images in every step, using the specifications from V1 and V2 could support approximately 3 simultaneous users, making scalability of the function difficult.

Crea.visions Version 4: Parisian Solutions: Participatory Design, Text-to-Image and an Online Universe

The main purpose of V4 is to provide support for civic problem solving (DG2, DG4), add meaningful short interactions (DG1,DG2), and to include text-to-image technology to expand creation possibilities (DG1).

Multi-stakeholder Alignment and Experimental Design

V4 was developed for the 2023 Learning Planet Festival, held in January in Paris, France. The events included three in person and two virtual workshops with bachelor and masters students at the Learning Planet Institute (LPI) and the general public. The events were advertised on both the Learning Planet Festival as well as Learning Planet Institute websites and social media. The new features and user flow was determined after holding two participatory design Sanders, Brandt, and Binder (2010) workshops (N=19) in November

2022 with Parisians. Participatory design is a method underscored by the CCS framework to engage community stakeholders in designing the tools and interventions as opposed to simply participating in or with a final version (DG4). Participation in all events was voluntary, but during image creation participants were provided with refreshments.

Technical Considerations

Similar to V3, the GAN has been retrained based on crowdsourced photographs from Paris and stock images to ensure diversity and representation of landmarks in Paris. V4 introduces an online companion app extending the features of the game (described below). To alleviate the load on the GPU node, in the companion app the blender is replaced with an interactive gallery of pre-generated Paris-GAN images with progressive navigation towards preferred features (Fig. 10). Additionally, due to the new open source availability of text-to-image generation, ***we designed the first workflow todate which combined image blending modalities of both sliders and text-to-image*** with PlaygroundAI (https://playgroundai.com/)

Gameplay and User Experience

Participants were introduced to the centeral concepts in human-AI co-creativity, V4 of crea.visions and then prompted to define a problem and identify where in Paris it is severe. Then they generated an image of the problem using sliders, then refined their image by feeding it as input to a text-to-image stable diffusion model (Borji, 2022) using PlaygroundAI. Participants then gave their image a title, and brainstormed solutions to the problem and identified which existing NGOs could get people engaged to help solve the problem. Participants could work individually or in pairs. Based on additional user testing from V3, the advanced mode was classified as confusing, adding little value to the user experience, thus we decided to remove it to improve the game flow (DG1). To extend its reach V4 introduces an online companion app allowing users to create similar submissions more rapidly than in the game (DG1). Users could locate their problem and/or solution on an interactive map and rank and discuss each other's submissions. The companion app lists online submissions alongside ingame submissions in a submission gallery that can be filtered by problem category.

Descriptive Statistics, Exploratory Analysis, and Lessons Learned

There were approximately 100 participants in the workshops and 34 submissions. See Fig. 12 for an example submission. Both images and log data from the blending process in crea.visions were collected, as well as all images generated in the text-to-image tool including prompts. As a whole, participants enjoyed the interactive aspect of the study and found the text-to-image tool in addition to the slider based crea.visions helpful in exploring and refining ideas. A survey was administered to the participants, which included questions about their prior experience with image blending tools and their overall experience with using the tools in a creative problem-solving setting. The results showed an equal distribution of answers between those who had prior experience with image blending tools and those who did not. The most popular topic that participants addressed with the tools was the environment. 50% of the participants reported being very interested in the problem they selected. Based on the survey, the tools were found to be helpful in visualizing and supporting idea exploration, as well as improving the quality of the results. For example, one participant commented, "Crea.visions made it easier to see how things would look like, if the problem was tackled properly." Weaknesses of the workshops included the lack of integration between the tools (e.g., participants had to move between crea.visions, Playground AI, and the companion app), the lengthy questionnaire, as well as the difficulty in prompt engineering.

Discussion and Future Work

The main novelty of this platform is its combination of human-AI co-creative technology with principles from the CCS framework to engage the public in image generation enabling civic expression and communication of societal challenges such as climate change. We approached the development of crea.visions with design goals of a user-friendly and intuitive interface (DG1), meaningful content and motivation for activities (DG2), meeting technical requirements (DG3), and community buy-in and co-design with multiple stakeholders (DG4). With respect to the usability improvements for exploring the vast GAN space (DG1, DG3), initial user exploration was supported through a single button generating new random images rather than manually adjusting each slider and base image. Furthermore, an experiment was conducted with dynamic visualization of the immediate possibility space around the current solutions by presenting 16 close-lying images. While this feature was useful to many participants, it was also computationally costly and distracting to others. As such, it may be reintroduced in future versions but was abandoned for now. Instead, the powerful text-to-image technology was explored for the V4 (DG1, DG3).

In terms of enriching the submission format, it was found that images could not stand alone, so they were supplemented with titles and short narratives (added in V2), which were often compelling, emotional and artistic. In V4, the submission format was enriched with geotagging, solution definition, NGO contact details and a chat feature. Exploratory quantitative and qualitative analysis of the submissions in each version have been performed, providing initial insights into the type of analysis that could be done in the future.

In coming iterations, the trend towards richer and smoother interactions as well as increased and lasting community impacts will be continued (DG2, DG4). This includes travelling exhibitions and using the platform as personal and approachable initiators of complex value-based discussions. A concrete proposed activity is having policy makers start their panel debates with generated images and their personal visions of the future. Concretely, in 2023, we

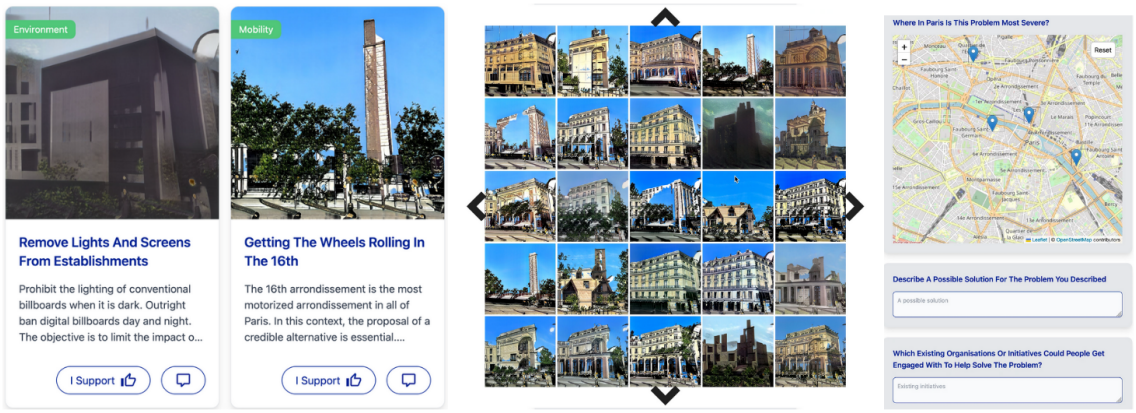

Figure 10: Screenshots from V4 companion app. Left: Submission Cards with tags from the submissions gallery. Middle: Interactive GAN image selector. Right: Interactive map and prompt fields in the submission form.

Figure 11: Example image created during the Paris workshops. Participant's Caption: The concept is a new way to move in a city. We don't use the roads because there are a lot of vehicles. That's a ecological transport and we choose the most used roads to take the most people possible..

are planning to run events in the Botanical Gardens in Belgrade, Serbia and at the Central Library, in Aarhus Denmark.

The potential of casual creators with a purpose goes beyond their immediate application in the crea.visions project. They can serve as powerful conversation starters and debate initiators across various settings and demographics, from children and the elderly to individuals from different languages and cultures. By integrating AI with human creativity and applying the CCS principles, our work indicates that one can facilitate inclusive and insightful discussions about societal and climate related topics which could also be useful in numerous fields, from education to policy-making.

Conclusion

In conclusion, this article presented the development and evolution of crea.visions for human-AI co-creation, specifically for community engagement and social good. Through four independent experiments, the platform was tested with over 1,000 participants and 25,000 generated images. The platform has evolved over the course of three years to become increasingly tailored for community-specific action, with the most recent version allowing users to geotag problems within a specific part of a city and submit solutions that are linked up with existing NGOs. The article emphasized that this platform is a step towards empowering the public to contribute to building and visualizing a collective vision of current societal issues and possibilities for our future world through the use of AI enhanced support tools. We hope that this platform inspires research institutions, local authorities and civil society organizations across the world to partner up to continue developing AI-tools for community empowerment, and strengthening the link between science and civil society.

Acknoweldgements

The authors would like to thank SHAPE, CircleU, Carlsberg Foundation and Novo Nordisk Foundation for their generous support of this work.

References

Baas, M.; De Dreu, C. K.; and Nijstad, B. A. 2008. A meta-analysis of 25 years of mood-creativity research: Hedonic tone, activation, or regulatory focus? Psychological bul-letin 134(6):779.

Belsky, S. 2020. Creativity will be key to competing against ai in the future workforce–here's how. In World Economic Forum. Retrieved, volume 12, 10–20.

Borji, A. 2019. Pros and cons of gan evaluation measures. *Computer Vision and Image Understanding* 179:41–65.

Borji, A. 2022. Generated faces in the wild: Quantitative comparison of stable diffusion, midjourney and dall-e 2. arXiv preprint arXiv:2210.00586.

Chang, J., and Ackerman, M. 2020. A climate change educational creator. In ICCC, 77–80.

Compton, K., and Mateas, M. 2015. Casual creators. In ICCC, 228–235.

Davis, N. M. 2013. Human-computer co-creativity: Blending human and computational creativity. In Ninth Arti-ficial Intelligence and Interactive Digital Entertainment Conference.

Epstein, Z.; Schroeder, H.; and Newman, D. 2022. When happy accidents spark creativity: Bringing collaborative speculation to life with generative ai. arXiv preprint arXiv:2206.00533.

Frich, J.; MacDonald Vermeulen, L.; Remy, C.; Biskjaer, M. M.; and Dalsgaard, P. 2019. Mapping the landscape of creativity support tools in hci. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 1–18.

Gajdacz, M.; Rafner, J.; Langsford, S.; Hjorth, A.; Bergenholtz, C.; Biskjaer, M. M.; Noy, L.; Risi, S.; and Sherson, J. F. 2021. Crea. blender: A gan based casual creator for creativity assessment. In ICCC, 405–409.

Hsu, Y.-C., and Nourbakhsh, I. 2020. When human-computer interaction meets community citizen science. Communications of the ACM 63(2):31–34.

Kantosalo, A., and Toivonen, H. 2016. Modes for creative human-computer collaboration: Alternating and task-divided co-creativity. In Proceedings of the Seventh International Conference on Computational Creativity, 77–84.

Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the im-age quality of stylegan. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 8110–8119.

Luccioni, A.; Schmidt, V.; Vardanyan, V.; and Bengio, Y.2021. Using artificial intelligence to visualize the impacts of climate change. IEEE Computer Graphics and Applications 41(1):8–14.

Rafner, J.; Hjorth, A.; Risi, S.; Philipsen, L.; Dumas, C.; Biskjær, M. M.; Noy, L.; Tyl´en, K.; Bergenholtz, C.; Lynch, J.; et al. 2020. Crea. blender: a neural network-based image generation game to assess creativity. In Extended abstracts of the 2020 annual symposium on computer-human interaction in play, 340–344.

Rafner, J.; Langsford, S.; Hjorth, A.; Gajdacz, M.; Philipsen, L.; Risi, S.; Simon, J.; and Sherson, J. 2021. Utopian or dystopian?: using a ml-assisted image gener-ation game to empower the general public to envision the future. In Creativity and Cognition, 1–5.

Rafner, J. 2021. Creativity assessment games and crowd-sourcing. In Creativity and Cognition, CC '21. New York, NY, USA: Association for Computing Machinery.

Sanders, E. B.-N.; Brandt, E.; and Binder, T. 2010. A framework for organizing the tools and techniques of partici-patory design. In Proceedings of the 11th Biennial Par-ticipatory Design Conference, PDC '10, 195–198. New York, NY, USA: Association for Computing Machinery.

Santanen, E. L.; Briggs; O, R.; Vreede; and De, G.-J. 2004 Causal relationships in creative problem solving: Com-paring facilitation interventions for ideation. Journal of management information systems 20(4):167–198. Schleicher, A. 2010. The case for 21st-century learning - OECD. Accessed: 2021-04-14.

Sielis, G. A.; Tzanavari, A.; and Papadopoulos, G. A. 2009. Enhancing the creativity process by adding context aware-ness in creativity support tools. In International Confer-ence on Universal Access in Human-Computer Interac-tion, 424–433. Springer.

Varkarakis, V.; Bazrafkan, S.; and Corcoran, P. 2020. Retraining stylegan-a first step towards building large, scal-able synthetic facial datasets. In 2020 31st Irish Signals and Systems Conference (ISSC), 1–6. IEEE.

Yannakakis, G. N.; Liapis, A.; and Alexopoulos, C. 2014.Mixed-initiative co-creativity..