Cat Royale: A Case Study of Artist-led AI Research

Dingbang Cao[1], Nick Tandavanitj[2], Matt Adams[2], Ju Row Farr[2], Eike Schneiders[3],
Guido Salimbeni[1], Mike Pound[1], Clara Mancini[4], Steve Benford[1]
[1]University of Nottingham, [2]Blast Theory,
[3]University of Southampton, [4]The Open University
steve.benford@nottingham.ac.uk

February 2025

Abstract: We explore artist-led research as a method to complement technical AI methodologies. We present a case study called Cat Royale in which artists created a robot to play with cats. We show how the artist-led development of this system involved extensive improvisation to create a socio-technical AI system that ultimately delivered a unique dataset of cats interacting with robots. We introduce a machine learning tool that enables diverse stakeholders to explore this dataset. We reflect on the distinctive characteristics of of artist-led AI research, the potential benefits to AI, and the tensions involved.

Keywords: Art, AI, Artistic-AI, Artistic Methods, AI Methodology, Dataset, Animal Behaviour.

1. Introduction

AI is transforming the arts and machine learning in particular has proven to be a powerful tool for contemporary artists [1] [2]. While there is increasing recognition of AI Art in the AI research community, including dedicated symposia and 'art tracks' at mainstream conferences, there has been less discussion of the converse relationship—*how can art transform AI*? We address this question by considering how artistic-practice can serve as a productive AI methodology that can complement existing data-driven technical methodologies. We do this through a case study of an artist-led exploration of AI (specifically robotics) that generated a groundbreaking artwork while also serving as a research project. Cat Royale was an artwork that explored the question of whether we should trust AI to care for our loved ones by creating a luxurious 'cat utopia' for a family of three cats, at the centre of which a robot tried to increase their happiness by playing games with them. A key outcome of the artwork was novel video dataset of cats interacting with robots which inspired the development of a tool to enable diverse researchers to better explore and analyse the data. We reflect on this case study to articulate the distinctive nature of artist-led AI research, identifying its potential benefits to AI research, while also noting key tensions that it raises.

2. Related Work

2.1. AI methodology

Methodology in AI broadly refers to the technical processes and methods used to create intelligent systems. There is no single AI methodology, but rather a collection of methods including machine learning, deep learning, natural language processing, computer vision, genetic algorithms, and evolutionary algorithms among others, each combining techniques from computer science, mathematics and statistics, often inspired by concepts from cognitive science and neuroscience. While different in their specifics, these methods can be broadly characterised as following a process of problem definition, data collection, data pre-processing, model selection and algorithm design, model training, evaluation and testing, deployment and verification, followed by continuous improvement. Underlying them are some important assumptions: that the problems we wish to solve are known in advance; that suitable data is available beforehand; and that there is an objective function representing a true and verifiable solution.

Critiques of AI's methodologies date back nearly as far as AI itself. The philosopher Hubert Dreyfus critiqued Classical AI for assuming that the human mind works by performing discrete computations on symbols, that all activity can be formalised mathematically, and that reality comprises mutually independent, atomic (indivisible) facts [3]. Dreyfus' argued that, to become truly intelligent, AI would need to be physically and socially embodied in the world, a view subsequently echoed in works such as Hutchins' theory of 'distributed cognition' [4] and Brooks' call for robotically embodied AI [5]. Even though AI's methodologies have since radically shifted to the data-driven approach of machine learning, these core assumptions remain in place. Moreover, new methodological critiques have emerged, notably ethical concerns around fairness and bias arising from inappropriate data [6], leading to calls to better document data provenance [7].

2.2. AI Art

We consider how art can extend AI's methodologies, potentially enabling them to respond to the above critiques. While the history of AI Art dates back to the earliest days of AI itself, the recent emergence of LLMs and text-to-image diffusion models has fuelled an explosion of interest in AI generated art through the widespread use of publicly available platforms for generating images (e.g., Google Imagine 3[1]), music (e.g., Suno)[2][3], and video (e.g., OpenAI Sora[4]). In turn, this has raised further ethical critiques of AI concerning intellectual property [8], job dislocation in the creative sector [9, 10], and the extent to which AI can be considered to be a creative artist or co-creator [11]. Various authors have discussed the relationships between artists and AI. Salimbeni et al. identified perspectives adopted by AI artists: employing AI as a tool for co-creativity, considering it an autonomous artist, creatively playing with data, and investigating and often critiquing AI as the subject of their work [11]. In a review of AI-generated visual art, Sivertsen et al. revealed how artists subvert AI's technical methodologies by deliberately selecting 'inappropriate' training data or overfitting or underfitting models so as to deliberately introduce ambiguity, thereby challenging widespread assumptions about the need for certainty, correctness and explainability [12]. Benford et al. consider how artists challenge the assumptions underpinning AI's technical methodologies, emphasising ambiguity above certainty, interpretation above explainability, improvisation above dependability, surrender above control, and playfulness above safety [13].

2.3. Artistic practice as a technology research method

Artistic practice is increasingly recognised as a productive research methodology that contrasts with established scientific methodologies. In the arts, practice-led research begins with creative artistic explorations, before reflecting on these to generate new knowledge, with research goals and questions often being unclear at the outset and only emerging through the artistic process [14]. Creative, practice-led research methodologies have migrated from the arts into in the field of Human-Computer Interaction (HCI) under the broad umbrella of Research Through Design [15, 16]. Of particular relevance here is the methodology of Performance-led Research in the Wild where professional artists create new artworks or installations, while technology researchers help them realise these technically, before reflecting on their artistic rationale and process and on audience engagement to distil new design principles for digital technologies [17].

[1] https://deepmind.google/technologies/imagen-3/

[2] https://suno.com

[3] Numerous high profile musicians have called for the acknowledgement of artists rights in relation to copyright and use of their creation for the training of AI models, e.g.,https://www.theguardian.com/music/2025/jan/27/elton-john-paul-mccartney-criticise-proposed-copyright-system-changes-ai

[4] https://openai.com/sora/

In what follows we present a case study of how Performance-led Research in the Wild worked as a method for conducting AI research, describing its process and outcomes before reflecting on the distinctive benefits and challenges that it might raise for AI research as a complement to more traditional scientific, data-driven and generally techno-centric methodologies.

3. Cat Royale: a case study of artist-led AI research

In August 2020, the UKRI funded the Trustworthy Autonomous Systems Hub (TAS Hub[5]), a UK research programme focused on investigating how autonomous systems can be created to be trustworthy by design. As part of this endeavour, the TAS Hub recognised that this goal required the involvement of diverse stakeholders including the public. To achieve this, the TAS Hub Creative Programme was born as an initiative aimed at using creative means of engaging the public and mainstream media. The artists Blast Theory were engaged as the TAS Hub's Creative Ambassadors and were invited to create a high profile work to engage the public with the question of trust in autonomous systems. Their response was Cat Royale [6].

3.1. An overview of Cat Royale

The artists created what they ambiguously called a 'cat utopia', a bespoke environment inhabited by a family of three cats intended to cater to their every need. Following the advice of animal welfare experts, the artists designed an enclosure containing ample feeding stations, litter trays, sleeping areas, perches and raised walkways at the centre of which was a robot arm whose task was to enrich the cats' lives by playing games with them. Every ten minutes, the robot would pick up and wield a toy, while the artists, observing from behind the scenes would score their response using the Participation in Play scale [18], a recognised instrument for measuring feline play, feeding the results back into a 'decision engine' allowing the artists, the robot, and its underlying systems to learn the cat's preferences and favourite games. Also located behind the scenes, a human robot operator would monitor that the robot was behaving safely and occasionally would step in to improvise movements when the robot became stuck, for example when it became tangled with the toys. The cats spent six hours a day for twelve days in the environment, engaging in over 500 playful games with the robot. They were continuously filmed via eight iPhones deployed at fixed positions, with the footage being edited into highlights for social media[7] and an eight-hour long film that has toured galleries worldwide, inviting audiences to consider whether they would trust a robot to care for their loved ones.

3.2. Cat Royale as a research project

Cat Royale was a research project as well as an artwork, following the method of Performance-led Research in the Wild described earlier. Initial articles from the project focused on how the work surfaced new requirements for designing *multispecies robot worlds*; environments designed to accommodate robots, humans and animals, in which AI and humans would collaborate to deliver care for the animals [19], alongside wider reflections on the complexities of ethical review processes for multispecies research [20]. These articles made contributions to the fields of Human-Computer Interaction and Animal-Computer Interaction. The question we consider here is the contribution to AI research.

The artists and wider research team had been keen to explore the potential of AI from the outset. Initial thoughts were to develop a computer vision system that could help monitor the cats' safety by predicting when they were approaching the robot (entering proximity or even crouching ready to pounce), or measuring how happily the cats were playing with the robot so that it could learn which games and movements they preferred, recommend new ones, or even generate new movements. However, given the lack of available training data on cats playing with robots, rigorous safety requirements, and the need to deploy a fully integrated and working system within a tight timescale, the final system required extensive human involvement to deliver these functions. A robot operator and a cat welfare officer monitored the cats' safety from a control room behind a one-way mirror. The robot operator continually pressed a dead

[5]https://tas.ac.uk
[6]Introductory video at: https://youtu.be/yBvHzsIgNRk
[7]https://www.youtube.com/watch?v=Un-bet9CFIQ&list=PLNEbg2a9DdkrlsKwPCq8JQDlekE7MGJrK

man's switch for the robot to proceed with its pre-recorded movements and sometimes needed to improve new movements to untangle the robot. The artists manually scored the cat's play, feeding the results into a simple rule-based system that recommended new games. They also manually created new movements in response to their observations. In short, while Cat Royale successfully deployed a physical robot to play with cats, the intelligence behind the robot was largely human, not artificial.

3.3. The Cat Royale dataset

However, a further key outcome of Cat Royale, which was not anticipated at the outset, was to deliver a rich dataset from which future AI systems might now learn.

Video data collection. Cat Royale yielded a unique dataset of cats engaged in a wide variety of activities, including playing with the robot, drinking, sleeping, grooming, scratching, and play fighting. The core of this was over 5TB of video data comprising 576 hours of footage captured from eight iPhone cameras at 30 fps and 1920*1080 resolution using the HEVC/H.265 codec. Seven cameras remained static throughout, while one camera angle was altered by the cats, and subsequently had to be repositioned. Camera positions (see Figure 1) were chosen to maximise coverage of the enclosure and capture the widest possible range of activities while providing overlapping views of key areas, especially the robot. Video data was supplemented by manually generated logs for each of the (over 500) games played, including a time-stamped record of which toy/game was deployed and the artists score of how the cats engaged using the Participation in Play Scale. This was further supplemented with frequent reports from the Cat Welfare officer using the Cat Stress Scale [21]. The wider data set also includes field observations (including additional video material recorded in the control room), interviews with key stakeholders, and documentation of the artists' development process including a 3D graphical—and physical—model of the enclosure.

Frame decimation and model training. We applied machine learning to develop a prototype tool for visualising and interacting with the data. This allowed the team to engage with the data and discuss it with key project stakeholders, thereby exploring potential uses and further developments of this dataset.

Following initial experimentation, we performed frame decimation to only sample one frame per second. This resulted in an acceptable trade-off between accuracy, time consumption, and storage requirements, resulting in a sampled data size of approximately 6 GB per 3-hour long video. Frame decimation to below .33fps, resulted in frequent missing of cats as they rapidly traversed camera views. Conversely, greater than 2fps significantly increased the storage space required and analysis time.

Following the frame decimation, we used a human-in-the-loop approach to dataset annotation, in which an existing object detection model was used to generate initial ground truth data, which was then corrected by a human annotator. We evaluated several popular models for object detection, including You Only Look Once (YOLO) [22], Single Shot Multibox Detector (SSD) [23] and Faster R-CNN [24]. We selected Faster-RCNN as our final model, which as a region-proposal-based model offers lower efficiency but improved accuracy on a small number of targets [25]. Evidence also suggests that Faster R-CNN offers higher performance on occluded animal targets [26], a common occurence in our dataset. While single-shot detectors such as YOLO are often more efficient, this was unnecessary in our experiments since annotation did not need to be conducted in real-time.

Cat Recognition and Regions of Interest (ROIs). To detect individual cats and Regions of Interest (ROIs) within the environment, we used the open source Detectron2 framework and applied a pre-trained Faster-RCNN model on randomly selected frames, before using LabelMe [27] to correct erroneous label suggestions. The model was then retrained with the corrected dataset before being used to recognise further images. This process was repeated until a total of 4320 correctly labelled frames were collected. Following the training, the model was applied across the larger dataset to identify the cats. Analysing a 3-hour video took between 2–3 hours on an Nvidia A4000 GPU.

We manually labelled static Regions of Interest (ROIs) within the environment, specifically: the robot, litter tray, water fountain, scratching post, grass, sleeping nest, right climber, left climber, walkway, and the door. To identify whether a specific cat was present in a given ROI at a given time, we analysed multiple

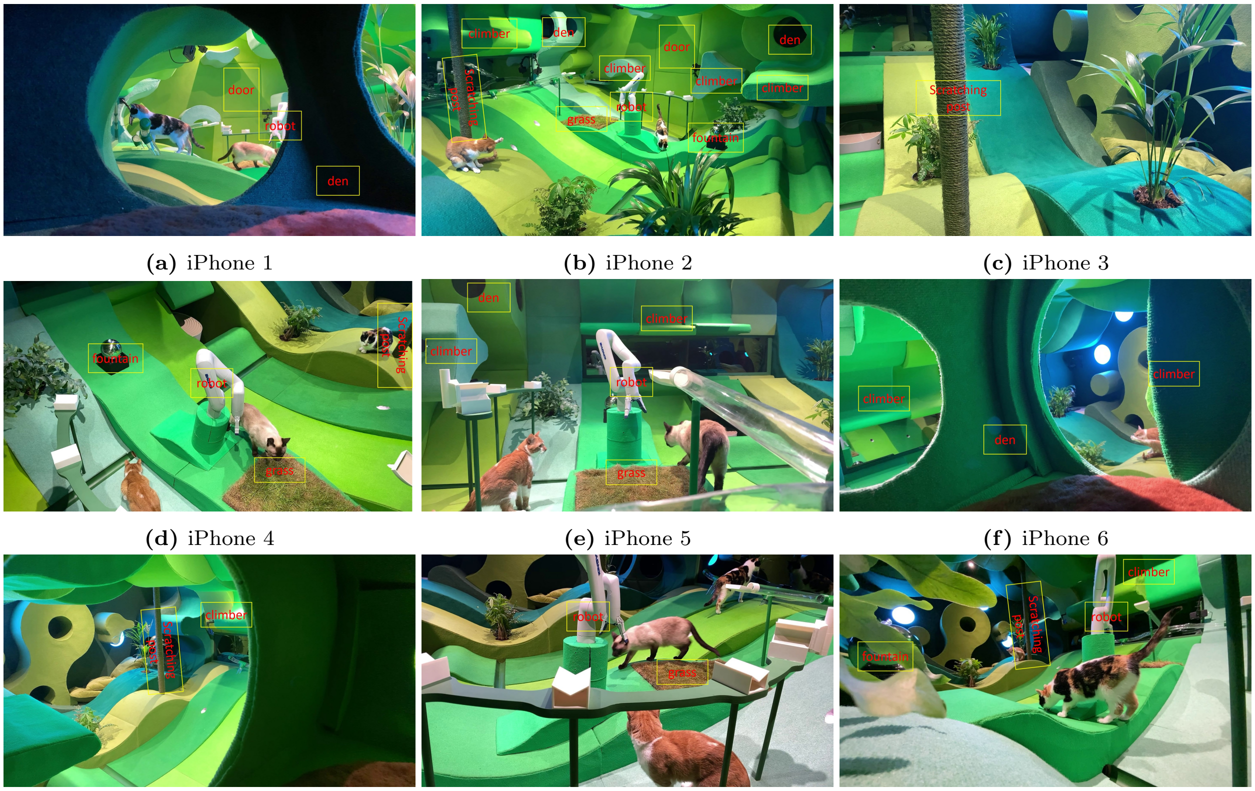

(a) iPhone 1 **(b)** iPhone 2 **(c)** iPhone 3

(d) iPhone 4 **(e)** iPhone 5 **(f)** iPhone 6

(g) iPhone 7 **(h)** iPhone 8 (perspective 8) **(i)** iPhone 8 (perspective 9)

Figure 1. Camera Perspectives. **1a:** Close-up perspective from inside from inside a floor-level sleeping den. **1b:** An overview of the main area with the robot, water fountain, climbing areas, scratching pools and plants visible. **1c:** Close-up of the scratching post, numerous plants, and the back wall of the enclosure. **1d:** Top-down view of the robot area. **1e:** Side view of the robot area, additionally this angle presents the entrance of the ball run as well as the rack for the placement of toys. **1f:** Close-up perspective from inside one of the high dens available to the cats. **1g:** Perspective from within another of the high dens available. This angle emphasises the scratching pole, several climbing areas, and the back wall of the enclosure. **1h:** Perspective of the toy rack, the robot, the ball run, and the central play area. **1i:** Following the bumping into the camera position 1h, this angle was chosen for camera 8. It focuses on the robot from a low angle.

camera views. If the cat was detected in more than half of the cameras covering that ROI—as we defined the certainty threshold to 50%—we considered it present in the region at that time.

Model evaluation. Our evaluation of the model was carried out on the Computer Vision Laboratory's GPU Cluster at the University of Nottingham using SLURM Workload Manager: CentOS (v. Linux release 7.9.2009), Conda (v. 23.10.0), Python v. 3.8.18, Detectron2 v. 0.6, OpenCV v 4.7.0, Torch v. 1.12.1, CUBA v. 11.7, GCC v. 10.2.0, GPU NVIDIA RTX A4000 with two GPU cores, and eight CPU cores. The average precision and recall for classifying the cats were 0.71 and 0.82, respectively[8]. Pumpkin's fawn colour made it the easiest to be identified against the green/blue background (lowest False Negatives). However, at times the toys were classified as Pumpkin (highest False Positives). For Ghostbuster, whose colour is creamy, we observed the opposite tendency.

Cat Identification through Video Analysis (CIVA). To enable various stakeholders in the project to better understand the dataset and explore future applications and challenges, we created a GUI-based interactive visualisation platform called CIVA (Cat Identification through Video Analysis) (see Figure 2). CIVA allows users to select combinations of cats, ROIs and timescales, and generates various visualisations in

[8]Precision and recall varied for the individual cats. The average precision for the identification of Ghostbuster had the highest average precision (0.75), while Pumpkin had the highest average recall (0.84).

response. An example visualisation of cat activity—focusing on Pumpkin's time spent at the Water fountain, the Robot, and the Grass ROIs—can be seen on the left side in Figure 4 (Pumpkin on the 20th of March). This visualisation highlights that Pumpkin spent significant time near the robot, compared to the other two ROIs chosen in this comparison (yellow line).

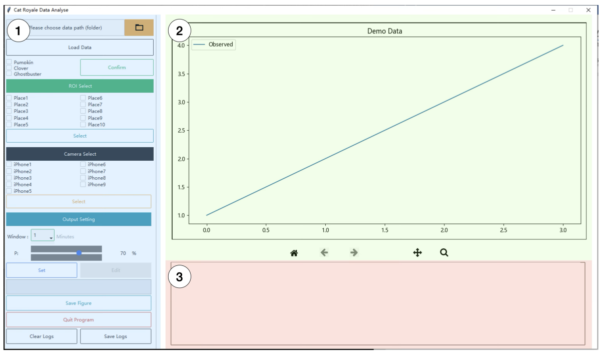

Figure 2. The CIVA system and its three main parts: 1) the area for parameter selection (blue), 2) the visual output (green), and 3) the data log (red).

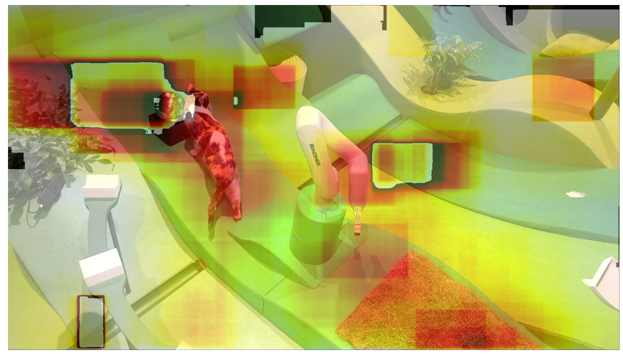

Figure 3. Heat map for Clover based on the perspective of Camera 4.

An alternative representation employs radar plots (see Figure 4 right) to compare the cumulative times spent by the various cats across of all ROIs. The final view currently supported by CIVA overlays heatmaps of cumulative time on video frames (see e.g., Figure 3). These example visualisations were not intended to present an exhaustive analysis of animal behaviour, but rather were intended to stimulate our stakeholders to imagine the kinds of analysis that might be possible in the future.

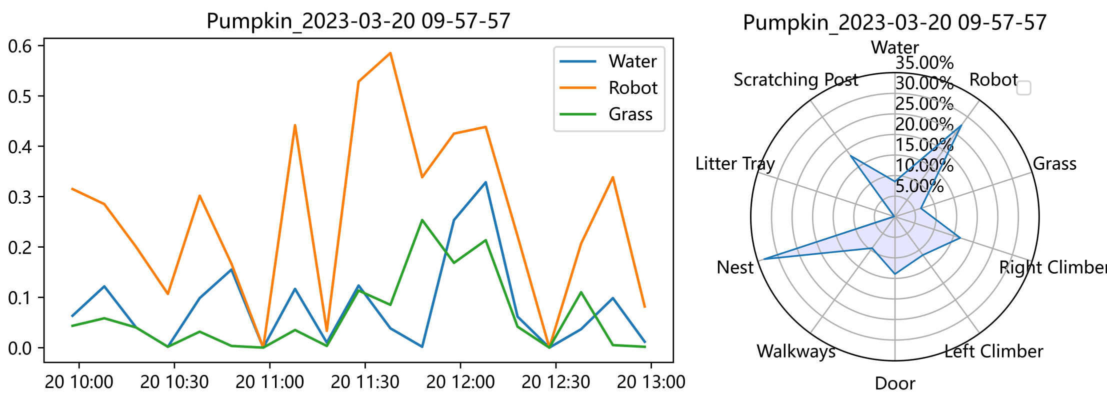

Figure 4. *Left:* An example representation of Pumpkins presence near a specific selection of ROIs, specifically: the Water fountain, the Robot, and the Grass. *Right:* Radar plot of how Pumpkin spend her time overall.

Initial responses to CIVA. We presented CIVA to five stakeholders of the Cat Royale project: one of the artists, an animal behavior researcher, an Animal-Computer Interaction (ACI) researcher, a computer scientist, and a roboticist —-interviewing them for an average of 45 minutes each to explore potential applications of the data set and identify potential applications and technical challenges that might arise. The artists, who had previously manually edited the Cat Royale movie and video highlights, anticipated how AI might help them better find stories of interest (e.g., all clips of Clover and the robot); might automatically edit the footage, especially if it could learn from how the human vision-mixer had done this; and might enable audiences to also better locate and watch key highlights. The two animal-focussed researchers suggested that the tool could summarise the cats' behaviours; help them identify key clips for subsequent manual analysis; and potentially be trained to automatically recognise specific cat behaviours such as playing, feeding, sleeping and scratching. They also speculated about the potential to develop future animal monitoring and care systems. The computer scientist and roboticist focused on how the dataset posed new challenges for AI video analysis and robotics, looking beyond current research in tracking animals and pose estimation to behaviour recognition arising from the datasets showcasing complex interactions, all consistently filmed from multiple angles.

4. Discussion: Artist-led AI research

We reflect on Cat Royale as a case study of artist-led research applied to AI, considering the distinctiveness of the approach, the potential benefits it offers, and the challenges it poses to AI research.

4.1. The nature of artist-led AI research

As noted earlier, the relationship between AI and art is not a new one; there have been many examples of employing AI to make art which in turn have demonstrated the capabilities of AI, raised questions and critiques, or served to shape the technology in some way. However, Cat Royale is distinctive in how it balanced making art with robotics and AI *research*, differing from conventional AI methodologies by being artist-led, while differing from many AI artworks in having an explicit goal to deliver research outputs, and involving researchers throughout. We highlight several important characteristics of Cat Royale that illustrate the nature of the approach.

The project was funded under a technology research programme and so was deeply embedded in AI research culture from the outset. However, it was artist-led, with Blast Theory being given free licence to determine the nature of the artwork. Their process was open ended and exploratory; the idea to work with cats emerged after six months while specific problems, objectives and research questions were not identified in advance beyond the broad goal of exploring public trust in autonomous systems. The artists adopted a nuanced stance towards AI, embracing the technology while inviting audiences to consider critical questions about its future.

In their turn, researchers played a dual role within the process, initially enabling the artists to realise their vision technically, but then to document and reflect on their rationale, process and the audience experience in order to distil new researcher findings for autonomous systems (e.g., the idea of designing multispecies robot worlds).

The method involved extensive improvisation of a socio-technical AI system. It was not possible to fully meet the artists' initial technical requirements given available data and time, so humans needed to step in to fill in the gaps. This shares similarities with the 'wizard of oz' methodology that is employed in robotics research [28], but at a scale and intensity that delivered a fully functioning experience sustained over many days. Finally, we note how research was multidisciplinary, drawing on and contributing to human-computer interaction, human-robot interaction, animal-computer interaction, and the humanities, as well as AI and robotics.

4.2. Potential benefits of artists' involvement to AI

We propose that this artist-led approach to the investigation of AI, its capabilities and limitations, can offer a breadth of potential benefits to AI research that complement established methodologies.

While research, technology development and academic writing—skills familiar to computer scientists, roboticist, animal-computer interaction specialists, and other academics—require creativity, the distinct creative expertise of artists introduces a unique and valuable perspective. Artists imagine unusual applications for AI, especially, emerging technologies, making their involvement invaluable. Collaborating with artists, provided they have artistic freedom, therefore provides a new and unique opportunity to investigate and test technology with unexpected applications and results.

Artistic practice often goes hand-in-hand with 'artistic thinking' that typically celebrates subjectivity, ambiguity, being open to multiple interpretations, improvisation and playfulness. Cat Royale's overall framing as a 'cat utopia' and the idea of measuring cats' 'happiness' were both highly ambiguous, defying simple 'correct' interpretations. Such thinking can broaden AI's conventional conceptual foundations that emphasise correctness, dependability and safety, providing responses to the critiques noted previously (see Section 2). The subjectivity inherent to art can highlight the perspectives and experiences of marginalised constituencies—animals in our case— drawing attention to fairness and bias and making the inherent positionalities (and political nature) of AI more apparent, especially to the public. This subjectivity extends to the CIVA tool that opens up the Cat Royale dataset to diverse human researchers to analyse video from different positions (e.g., animal behaviour, robotics etc).

The resulting artworks may deliver extensive public engagement as they continue to tour once initial production and research has finished. Researchers may benefit from powerful impact stories backed up by evidence of audience numbers, venues, additional funding through commissions, awards and critical acclaim. In turn, artists may benefit from co-authorship of research outputs which may ultimately enable them to act as independent research organisations in projects (Blast Theory appear as co-authors of technical papers from Cat Royale and have been funded in several EU and Innovate UK projects).

Importantly, AI artworks delivered at scale can yield rich and unusual new datasets for the wider AI community. This reverses the conventional approach to AI in which the availability of data precedes the development of systems and then experimental validation. In Cat Royale, the absence of suitable data at the outset necessitated the improvisation of systems, which ultimately delivered new data. Moreover, capturing rich documentation of artworks is a widely researched topic in the humanities [29], with the implication that humanities scholars may have an important role to play in extending current approaches to capturing and structuring artistic datasets for AI.

4.3. Tensions in artist-led AI research

It is important to recognise that artist-led research also raises tensions. The adventure of artist-led research brings a degree of risk, with success depending on public delivery of working performances to real audiences according to strict deadlines; a degree of rigour that far exceeds conventional technology demos. Researchers encounter further risks arising from touring, including securing opportunities to perform and managing all of the complexities of touring a professional artwork.

There is a clear epistemological tension between science and engineering, on which AI conventionally draws, and the arts and humanities. With the artist-led approach, research questions may only emerge during the process while evaluations may be qualitative and highly subjective. Consequently, it can be difficult to persuade the 'scientific' community to recognise the validity and value of research contributions, making it necessary to clearly and explicitly clarify the methodology.

The risks of the approach may extend to ethical review which needs to consider ethical issues arising for engaging audiences with novel and sometimes provocative experiences. The distinctive ethical considerations of artist-led research have been articulated in [30], while Cat Royale was notable for its eighteen month long ethical review process (in large part because it also involved animals) which had to negotiate various epistemological tensions as reported in [20].

Finally, it may be difficult to balance funding for projects. AI research funders may be tempted to view the approach as 'public engagement' more than research (as was initially the case for Cat Royale), while it is also important for artists to secure commissions and arts funding for production and touring, in part to demonstrate the artistic legitimacy of the work. A final important point here concerns payment for the artists. Research funders often view industry partners as providing (typically match) funding to a project rather than being paid. However, as small independent companies, it is important that artists are fairly remunerated for participating in research.

5. Conclusion

The case study of Cat Royale reveals how artist-led research can be a distinctive and productive research method for AI. It's distinctiveness lies in enabling artists to pursue an exploratory and improvisational approach to delivering functioning AI systems that work 'in the wild' of public performance. Unlike conventional machine learning models in which the availability of data precedes models, this approach delivers unusual datasets as an output of the project that may then raise new challenges for the wider community. There are also tensions inherent to this approach, the most immediate of which lies in recognising it as a legitimate research method despite its epistemological differences to conventional techno-centric ones.

References

[1] Anne Ploin, Rebecca Eynon, Isis Hjorth, and Michael A Osborne. Ai and the arts: how machine learning is changing artistic work. report from the creative algorithmic intelligence research project, 2022.

[2] Christian Then, Erickar Soewandi, Muhammad Danial, Said Achmad, and Rhio Sutoyo. The impact of artificial intelligence on art - a systematic literature review. pages 1–7, 10 2023.

[3] Hubert L Dreyfus. Alchemy and artificial intelligence, 1965.

[4] Edwin Hutchins. *Cognition in the wild*. MIT press, 1995.

[5] Rodney A Brooks. Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159, 1991.

[6] Kate Crawford. The atlas of ai: Power, politics, and the planetary costs of artificial intelligence, 2021.

[7] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, 2021.

[8] Simon Chesterman. Good models borrow, great models steal: intellectual property rights and generative ai. *Policy and Society*, page puae006, 02 2024.

[9] Harry H. Jiang, Lauren Brown, Jessica Cheng, Mehtab Khan, Abhishek Gupta, Deja Workman, Alex Hanna, Johnathan Flowers, and Timnit Gebru. Ai art and its impact on artists. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, page 363–374, New York, NY, USA, 2023. Association for Computing Machinery.

[10] Reishiro Kawakami and Sukrit Venkatagiri. The impact of generative ai on artists. In *Proceedings of the 16th Conference on Creativity & Cognition*, C&C '24, page 79–82, New York, NY, USA, 2024. Association for Computing Machinery.

[11] Guido Salimbeni, Steve Benford, Stuart Reeves, and Sarah Martindale. Decoding ai in contemporary art: A five-trope classification for understanding and categorization. *Leonardo*, pages 420–426, 2024.

[12] Christian Sivertsen, Guido Salimbeni, Anders Sundnes Løvlie, Steven David Benford, and Jichen Zhu. Machine learning processes as sources of ambiguity: Insights from ai art. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2024.

[13] Steve Benford, Adrian Hazzard, Craig Vear, Helena Webb, Alan Chamberlain, Chris Greenhalgh, Richard Ramchurn, and Joe Marshall. Five provocations for a more creative tas. In *Proceedings of the First International Symposium on Trustworthy Autonomous Systems*, pages 1–10, 2023.

[14] Linda Candy, Ernest Edmonds, and Craig Vear. Practice-based research. In *The Routledge international handbook of practice-based research*, pages 27–41. Routledge, 2021.

[15] John Zimmerman, Jodi Forlizzi, and Shelley Evenson. Research through design as a method for interaction design research in hci. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 493–502, 2007.

[16] William Gaver. What should we expect from research through design? In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 937–946, 2012.

[17] Steve Benford, Chris Greenhalgh, Andy Crabtree, Martin Flintham, Brendan Walker, Joe Marshall, Boriana Koleva, Stefan Rennick Egglestone, Gabriella Giannachi, Matt Adams, et al. Performance-led research in the wild. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 20(3):1–22, 2013.

[18] Jacklyn J. Ellis. Beyond "doing better": Ordinal rating scales to monitor behavioural indicators of well-being in cats. *Animals*, 12(21), 2022.

[19] Eike Schneiders, Steve Benford, Alan Chamberlain, Clara Mancini, Simon Castle-Green, Victor Ngo, Ju Row Farr, Matt Adams, Nick Tandavanitj, and Joel Fischer. Designing multispecies worlds for robots, cats, and humans. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 2024.

[20] Steven David Benford, Clara Mancini, Alan Chamberlain, Eike Schneiders, Simon D Castle-Green, Joel E Fischer, Ayse Kucukyilmaz, Guido Salimbeni, Victor Zhi Heung Ngo, Pepita Barnard, Matt Adams, Nick Tandavanitj, and Ju Row Farr. Charting ethical tensions in multispecies technology research through beneficiary-epistemology space. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 2024.

[21] M R Kessler and D C Turner. Stress and adaptation of cats (felis silvestris catus) housed singly, in pairs and in groups in boarding catteries. *Animal Welfare*, 6(3):243–254, 1997.

[22] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.

[23] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 21–37. Springer, 2016.

[24] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2016.

[25] Shrey Srivastava, Amit Vishvas Divekar, Chandu Anilkumar, Ishika Naik, Ved Kulkarni, and V. Pattabiraman. Comparative analysis of deep learning image detection algorithms. *Journal of Big Data*, 8, 2021.

[26] Ivan Roy S. Evangelista, Lenmar T. Catajay, Maria Gemel B. Palconit, Mary Grace Ann C. Bautista, Ronnie S. Concepcion, Edwin Sybingco, Argel A. Bandala, and Elmer P. Dadios. Detection of japanese quails (coturnix japonica) in poultry farms using yolov5 and detectron2 faster r-cnn. *Journal of Advanced Computational Intelligence and Intelligent Informatics*, 26, 2022.

[27] Bryan C. Russell, Antonio Torralba, Kevin P. Murphy, and William T. Freeman. Labelme: A database and web-based tool for image annotation. *International Journal of Computer Vision*, 77, 2008.

[28] David Sirkin, Brian Mok, Stephen Yang, and Wendy Ju. Mechanical ottoman: How robotic furniture offers and withdraws support. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, HRI '15, page 11–18, New York, NY, USA, 2015. Association for Computing Machinery.

[29] Gabriella Giannachi and Jonah Westerman. *Histories of performance documentation*. Taylor & Francis, 2017.

[30] Steve Benford, Chris Greenhalgh, Bob Anderson, Rachel Jacobs, Mike Golembewski, Marina Jirotka, Bernd Carsten Stahl, Job Timmermans, Gabriella Giannachi, Matt Adams, et al. The ethical implications of hci's turn to the cultural. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 22(5):1–37, 2015.