Beyond the Spotlight: Co-Designing AI for Theatre Audience Communication

Alan Pedrassoli Chitayat
University of York
York, United Kingdom
alan.pchitayat@york.ac.uk

Steph Carter
Department of Computer Science
University of York
York, United Kingdom
steph.carter@york.ac.uk

Barry Alan Robertson
The Grey Hill
Loans, United Kingdom
barry@thegreyhill.com

Jonathan Hook
School of Arts and Creative Technologies
University of York
York, United Kingdom
jonathan.hook@york.ac.uk

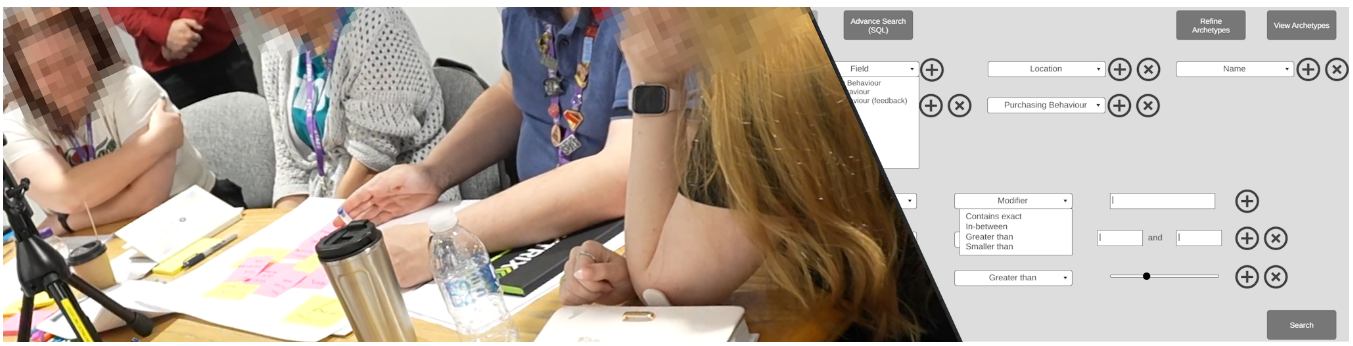

Abstract

Theatres and concert halls play a crucial role within the performing arts, where managerial and administrative staff are essential to bringing live performances to audiences. Existing AI research has focused on artistic creation, but less attention has been paid to the purposeful design of AI systems that support organisational practices. This paper addresses this gap by identifying the needs, challenges and opportunities for AI integration into everyday workflows, forming the basis for design principles to guide the architecturing, training, and deployment of AI systems that empower staff, rather than replace them. This is explored through a co-design workshop with theatre marketing and communication professionals. Through reflections of the themes explored in the workshop and by following the guiding principles, this paper presents examples of implementation of AI systems that could be adopted, offering concrete directions for developing AI that benefits the cultural sector.

CHI '26, Barcelona, Spain

ACM ISBN 979-8-4007-2278-3/26/04
https://doi.org/10.1145/3772318.3791710

CCS Concepts

• **Human-centered computing** → *Human computer interaction (HCI)*; *User studies.*

Keywords

Design, RTA, Theatre, AI, Co-design Workshop

ACM Reference Format:
Alan Pedrassoli Chitayat, Steph Carter, Barry Alan Robertson, and Jonathan Hook. 2026. Beyond the Spotlight: Co-Designing AI for Theatre Audience Communication. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3772318.3791710

1 Introduction

Theatre is a vital part of culture and and cultural heritage [2, 6, 12, 21]. Smaller theatres in particular rely on continuous communication with local audiences and donors, yet they operate with limited staff capacity and increasingly complex digital ecosystems [36, 49]. Failure to adapt risks the gradual loss of an important cultural domain [36].

Across the wider non-profit and cultural sector, digital communication and customer-relationship management have become central to organisational sustainability. Commercial marketing teams routinely draw on sophisticated CRM platforms, automated communication pipelines, and AI-supported analytic tools. However, most small cultural organisations - particularly local theatres - lack the resources, training, or data infrastructure required to implement such technologies at scale. Even where theatres use established CRM systems (e.g., Spektrix, Tessitura, Eventbrite), the existing AI-enabled features focus primarily on ticketing operations or high-level analytics, offering limited support for the day-to-day communication and marketing practices that staff must manage manually. At the same time, AI continues to reshape how people create, discover, and consume media [1, 9, 11, 22, 30, 32, 47, 51]. Despite these developments, relatively little attention has been given to how AI could support the organisational practices of theatres, such as communication and marketing. This mismatch raises a broader question of *how might AI systems be designed to meet the organisational realities, professional values, and sector-specific constraints of theatre communications work?*

This paper addresses that gap by examining the values, challenges, and practices of theatre marketing and communications professionals, and by exploring how AI could meaningfully support their work. Using a co-design workshop with eight staff members from a producing theatre, we applied Value Sensitive Design [13, 14] and Reflexive Thematic Analysis (RTA) [4] to surface how staff articulate their responsibilities, constraints, and aspirations for AI augmentation. While the study is intentionally focused on a single organisation, this choice enables in-depth engagement with an otherwise underexplored user group, we discuss how the findings align with challenges across comparable non-profit cultural institutions, and where transferability or divergence may arise.

The contribution of this paper is threefold. First, it presents novel insights into the values and needs of theatre marketing and communications professionals, offering one of the first systematic accounts of how people in these roles envision the use of AI. Second, it proposes a set of guiding principles for the design of AI systems in this underexplored context, providing actionable considerations for both researchers and practitioners. Third, it illustrates how these principles can be operationalised through two connected examples of system architectures, contributing concrete starting points for the development of AI that supports cultural institutions. Together, these contributions extend existing work on AI in the arts by shifting the focus from artistic creation to the organisational practices that sustain cultural heritage.

2 Related Work

2.1 AI in the Arts and the Cultural Sector

Research exploring AI in the arts has primarily focused on creative production and audience experience, rather than on the organisational practices that sustain cultural institutions. Recent examples include AI-assisted content generation across diverse artistic practices [15, 35, 44, 52], including systems such as Dramatron [23], which supports scriptwriting through large language models. Parallel work investigates how AI and machine learning can be used to evaluate artistic experience, for example through physiological or behavioural measurements of audience synchrony, immersion, and engagement [17, 17, 27, 42, 43, 50].

These strands demonstrate the increasing presence of AI in artistic domains, yet both centre artistic outputs rather than the professional labour involved in promoting and sustaining them. This is noteworthy given that cultural organisations, including theatres, now rely heavily on digital communication, audience engagement strategies, and donor relationships for financial resilience. While commercial marketing sectors have rapidly adopted AI-enabled customer analysis, automation, and content generation, similar systems are not readily accessible or appropriately tailored to small cultural organisations.

As a result, the organisational practices that enable audiences to discover, understand, and attend cultural events remain largely underexplored in HCI and AI in the arts research. Understanding the values and challenges of marketing and communication staff is, therefore, essential for designing AI that is meaningful in its support of cultural organisations beyond artistic outputs.

2.2 Value-Centred Design of AI Applications

Value Sensitive Design (VSD) offers an established approach for developing technologies that account for stakeholder values throughout the design process [14]. Its emphasis on contextual and ethical considerations aligns closely with user-centred design (UCD) methods [20], yet collaboration between AI practitioners and VSD or UCD researchers is not always realised in practice [24]. This disconnect can lead to data-centric AI systems that place the burden of adaptation on users rather than supporting their existing practices [38].

Recent work has sought to bridge this gap by embedding stakeholder needs into AI system development. For example, Sadek et al. [34] provide a toolkit for incorporating user values into AI design processes, while Pedrassoli Chitayat et al. [28] demonstrate how early exploration of user needs influences the effectiveness and acceptance of machine learning solutions. Across this literature, a recurring theme is the need for AI systems that are interpretable, transparent, and usable by professionals who may lack technical expertise.

For example, in marketing and advertising, AI is widely used for prediction, segmentation, and automated content creation [7, 8, 29, 46]. While such systems can offer powerful analytic capabilities, they frequently present challenges around interpretability, ethical implications, and the translation of model outputs into actionable workflows [19, 31]. These tensions become particularly salient in cultural contexts, where communication practices are shaped not only by commercial considerations but by social, cultural, and community-oriented values.

Consequently, there is a growing need for AI systems that respect the organisational realities of cultural institutions, support professional judgement, and align with the values they aim to promote.

2.3 Related Methodologies

Co-design and qualitative methods are increasingly used in the development of cultural and creative technologies, offering a means to foreground stakeholder values in complex and highly contextual domains. For example, Tzanidou et al. [45] used co-design with actors and acting students to develop a toolkit for accessible theatre experiences, drawing on Reflexive Thematic Analysis (RTA) [4] to identify stakeholder priorities and translate them into design principles. Similarly, Trajkova et al. [41] applied RTA to understand dancers' perceptions of robotic partners, using stakeholder values to guide the design of interactive AI systems.

These studies illustrate the utility of co-design and RTA for investigating user perspectives to guide the design of future technologies. Building on this tradition, the present work applies co-design and RTA within the organisational context of theatre marketing. Rather than focusing on creative or performance-oriented technologies, this study examines the everyday labour of communication professionals and uses their articulated values to derive guiding principles for AI systems that support audience engagement, donor relations, and organisational sustainability.

2.4 Customer Relationship Management Systems in Theatre Contexts

Customer Relationship Management (CRM) systems are widely used across the cultural and creative sectors to support administrative, financial, and audience management activities. Within theatre organisations, these platforms typically function as centralised systems for ticketing, donor management, audience records, and revenue reporting. Examples commonly adopted in practice include specialist arts CRMs such as Spektrix,[1] Tessitura,[2] and PatronManager,[3] as well as more general-purpose platforms adapted for arts use, such as Salesforce-based solutions.[4]

Prior research has highlighted the role of CRMs in supporting organisational efficiency and financial sustainability within cultural institutions. Studies examining digital infrastructure in the arts describe how CRM platforms enable organisations to consolidate audience data, manage transactions, and evaluate attendance patterns at scale [36, 49]. In this capacity, CRMs play an important role in stabilising operational workflows and supporting strategic planning, particularly for smaller organisations with limited resources.

Contemporary CRM systems increasingly incorporate data analytics and automation features. These include tools for audience segmentation, performance dashboards, and, in some cases, predictive models intended to support revenue forecasting or donor identification [7, 29]. Such capabilities align with broader trends in data-driven marketing, where automation is leveraged to improve efficiency and optimise decision-making. However, these features are primarily oriented towards aggregate analysis and operational outcomes, rather than the situated, interpretive work involved in crafting and managing day-to-day communications with audiences.

Importantly, while CRMs may include email campaign tools or integrations with external marketing platforms, they generally provide limited support for the creative and relational aspects of communication work. The design focus of these systems remains centred on ticketing operations, customer records, and transactional histories, rather than on assisting staff in deciding how, when, and in what tone to communicate with diverse audiences across multiple channels. As a result, theatre marketing and communications professionals often rely on additional tools and manual processes to translate CRM data into communicative practice.

[1]https://www.spektrix.com
[2]https://www.tessitura.com
[3]https://www.patronmanager.com
[4]https://www.salesforce.org

3 Methodology

This study investigates how AI-based solutions might be designed to support theatre staff in audience communications and marketing work. To explore this question, a co-design workshop was conducted with the marketing and communications team of Perth Theatre and Concert Hall. Focusing on the entirety of a single professional team enabled in-depth engagement with the shared practices, interdependencies, and organisational constraints that shape their daily work. While this approach does not aim to produce statistically generalisable insights, it allows for a detailed understanding of a coherent organisational unit whose responsibilities and challenges are representative of many small and mid-sized cultural organisations across the United Kingdom. Prior to data collection, a full description of the project was provided to the Ethics Committee of the University of York and approval of the project was granted.

The workshop data was analysed using Reflexive Thematic Analysis (RTA) [4, 5] to identify the values, challenges, and expectations that staff associate with potential uses of AI. Eight participants (P1–P8) took part, representing the full communications and marketing team, which includes staff responsible for campaign design, content creation, donor relations, digital operations, and audience engagement. Because the participants formed an existing team rather than a recruited sample, demographic data were not collected. The workshop lasted three hours and consisted of two structured activities: (1) identifying current challenges and pain points, and (2) developing user journeys that incorporated hypothetical AI solutions. Participants were divided into two groups of four (Group A: P5–P8, Group B: P1–P4), assigned randomly and kept consistent across tasks to enable sustained discussion. At the end of each task, groups presented and discussed their work with one another, to preserve the benefits of the collective team.

3.1 Co-design Workshop Task 1

The first task aimed to identify common challenges in participants' day-to-day work. Participants wrote their routine tasks on large sticky notes and then added smaller notes describing barriers, pain points, or difficulties associated with each task. This activity was designed to help participants situate themselves in the workshop by reflecting on their own practices.

Groups were given 15 minutes to list as many tasks and challenges as possible. They were then asked to select one task and its associated pain points to present to the other group, followed by 15 minutes of discussion. This discussion focused on understanding why these pain points were problematic and what consequences they created. At this stage participants were not asked to propose solutions. Open-ended questions from researchers were used to stimulate dialogue and encourage reflection.

3.2 Co-design Workshop Task 2

The second task introduced participants to user journeys [18] as a design tool. Researchers first illustrated the concept with an example journey of a patient registering at a local doctor's practice. This example highlighted steps such as "examine" (researching practices), "elect" (choosing one), "evaluate" (attending an appointment), and "entrust" (booking future visits). Sketches of the journey were used to show how visualisation can support both reflection and conversations while designing and presenting the journey.

After a short break, participants were asked to construct their own user journeys, focusing on a task from their professional practice. They were instructed to incorporate a hypothetical AI-powered tool into the journey, disregarding feasibility in order to focus on ideal functionality. The task lasted 45 minutes, during which researchers provided clarification and support if required. It is important to note that, as participants are not experienced designers, the quality of the user journeys themselves (and the adherence to conventional methodological steps) [33, 40] were not a priority during the workshop. Instead, the user journeys were put in place as a means to stimulate conversations, and ensure participants were considering a concrete use case.

Following the design activity, participants presented and discussed their user journeys for one hour. The discussion explored how AI systems could be integrated into workflows, how they might fit into daily routines, and what outputs would be most valuable. Concerns regarding data availability and interpretability were also raised, alongside suggestions for innovative data use. The results of these discussions are presented in Section 4 and reflected upon in Section 5.

3.3 Reflexive Thematic Analysis

Audio from the workshop discussions was recorded and transcribed. Analysis of the transcriptions focused on the two discussion periods in Tasks 1 and 2, and supplemented by the artifacts produced during the workshop, which is made available in Appendix A. RTA was conducted following established practice [5].

Codes were generated inductively from the data and then developed into themes. Participant artefacts were also consulted to support interpretation. Coding and theme development were carried out by the primary researcher, which were then presented to the second researcher for peer debriefing. This was done to ensure themes captured the data accurately. The researcher's positionality as an AI researcher with expertise in user-centred design was acknowledged as shaping both interpretation and theme generation.

The aim of the analysis was to identify guiding principles for incorporating AI into theatre marketing and communications. These principles, along with example implementations, are presented in Section 5.

4 Results

Through analysing the data collected during the co-design workshop, five major themes were identified. These themes are presented in detail in this section, including theme definitions and relevant quotations. Any participant references have been anonymised using P1–P8. Artifacts produced by participants are also provided - for reference - in Appendix A.

4.1 Workload Constraints Limit Flexibility

This theme is divided into two subthemes. Within both subthemes, participants describe how their current day-to-day activities offer them little-to-no flexibility to partake on any tasks, activities or responsibilities beyond their existing workload. Importantly, this does not reflect on a willingness to change their approach to performing any task, but instead depicts how they feel like they do not have the capacity to take on any additional responsibilities beyond what they are currently already doing.

4.1.1 *Challenges in Creating Content.*

Participants described a range of challenges associated with generating marketing and communication content. These challenges relate not to identifying targets or recipients, but to the work of content creation itself.

The diverse nature of the theatre's programme often makes it difficult to produce appropriate materials, especially when staff have limited familiarity with specific shows. This difficulty is compounded by a lack of clear, digestible information and by the fast-moving nature of the theatre's events. Participants noted that creating "*enough stuff to be out on it to create interest constantly*" (P5) was a persistent struggle.

Another challenge was the need to avoid errors when manually inputting variable information such as start times or run dates. Participants described relying on several rounds of internal checking to prevent mistakes. Alongside accuracy, they emphasised the importance of maintaining a human and engaging tone, with P4 explaining that they do not want communication to be "*just this show, this show, this show [...] We need some human [...] some story or like some other way of creatively showing your personality as a business.*"

4.1.2 *Time-Consuming Daily Tasks.*

Participants reported that their daily tasks leave them with no significant spare capacity. Routine work often takes longer than expected, making it difficult to estimate or schedule tasks. For example, writing an email was described as taking "*a good hour just to sort out one email*" due to the repetitive nature of the work (P6). Similarly, P3 noted that "*something that should be simple, can take a long time.*"

Another source of time pressure was the constant change in digital advertising platforms. Participants explained that they may spend several weeks learning how to prepare ads on Meta only for requirements to change, forcing them to re-learn established processes.

Participants also expressed a desire to analyse their data in more depth but noted that time constraints limit them to superficial analysis. As P6 put it, they only have "*a generalised concept of analysing the data*" because of "*the time that we have to spend doing other firefighting.*"

When imagining an ideal AI system, participants emphasised ease of use, minimal learning burden, and digestible outputs (such as simple charts rather than complex analytical tools). They also expressed interest in using AI to streamline existing processes so that they could dedicate more time to strategic work.

4.2 Granularity of Audience Understanding

This theme contains two subthemes that describe participants' desire to gain deeper insight into both broad audience patterns and individual audience members.

4.2.1 Identifying Audience Trends.

Participants expressed a need to better understand high-level audience trends. They highlighted the importance of segmentation by demographic factors such as age or background, noting that some events (e.g., certain tribute acts) appeal to "*different people from different backgrounds, different age groups*" (P3).

They also emphasised the need for continuous re-segmentation as preferences evolve and more data becomes available, using sources such as post-attendance feedback, purchase histories, and travel distance. Genre was identified as an important but underutilised variable, with participants noting that audience preferences for comedy, musicals, or other genres could meaningfully support segmentation (P2).

Participants were also interested in distinguishing between regular audience members and donors, and in identifying potential donors within their data. As P5 explained, it would be useful to identify "*which type of person is a donor*" and "*who else within our audience is a potential donor.*"

4.2.2 Understanding Individuals Personally.

Beyond aggregate trends, participants expressed a desire to understand individuals on a more personal level, especially donors who often develop personal relationships with staff. When describing an ideal system, P6 suggested that if a conversation at a donor event revealed personal details, such as "*a daughter's name*" for instance, the system could recall this during future interactions to serve as points to start conversations.

Participants were also interested in knowing when individual audience members preferred not to be contacted, or when repeat communications might be helpful, acknowledging that individuals vary significantly in their preferences and responsiveness.

4.3 Effective Audience Interaction

Participants outlined how they feel it is important to interact with audiences and customers appropriately and efficiently. This theme is composed of three subthemes, which deal with how the theatre (and its staff) interact with audiences, prospective audiences and donors.

4.3.1 Optimising Communication Channels.

Participants expressed concern about over-communicating. As P7 put it, "*we send so many emails, we don't want to inundate them with something we've already*" said. They also noted that different audience members prefer different channels—email, letters, phone calls, SMS, targeted ads, or social media posts.

They highlighted the importance of reciprocal communication and timely requests for feedback. P6 noted that some questions "*could be asked technically before the show but are better asked after*" , emphasising the importance of the timing of communication.

4.3.2 Targeting Without Excluding.

Participants described the challenge of deciding what to say, and to whom, without overwhelming or excluding audiences. One key concern was the signal-to-noise ratio in communications. P2 explained that they aim to avoid sending so much information that audiences "*automatically ignore it.*"

Participants also worried about overly narrow targeting practices. For example, shows with specific themes may appeal to a primary group (e.g., football fans) while also attracting broader audiences. P6 described the need to "*tailor different streams*" so as not to exclude others.

Participants also discussed creating multiple versions of communications to reach diverse groups. They expressed interest in using AI to identify relevant keywords, behavioural patterns, or geographical factors such as travel distance. P4 noted that tools able to surface commonly used search terms would help reach prospective audiences more effectively.

4.3.3 Ensuring Donor Benefits are Meaningful.

Participants stressed the importance of ensuring that donor benefits—particularly within the theatre's "friends" scheme—remain meaningful. They expressed interest in benchmarking their offer against other organisations, for example by examining "*the type of pricing [...] the discounts that they offer*" (P3). They also wished to understand which benefits donors "*actually want or have interacted with the most*" (P4) to ensure relevance and value.

4.4 Opportunities for Data-Driven Insights

This theme is composed of two subthemes, both of which describe participants desire to be more effective with the way their data is collected, analysed or otherwise utilised for their benefit.

4.4.1 Improving Data Clarity.

Participants frequently noted the lack of a centralised view of their data, describing current systems as fragmented. P1 summarised this challenge: "*We wanted to be able to interact with data across multiple platforms [...] so we know everything from what that individual is doing.*"

Genre information was again highlighted as inconsistently captured, making it harder to understand how it shapes audience behaviour. Participants also discussed "unknown unknowns" in their datasets, expressing a desire for AI systems that could identify missing fields or inconsistencies. For example, P6 described the need for tools that could highlight missing genre data.

Participants also expressed concern about personal bias in interpreting data. As P3 put it, "*there's a danger that you can based on your own assumptions, profile someone,*" when the data might indicate something different.

4.4.2 Leveraging Data Outputs.

Participants described interest in generating new insights that they currently cannot produce. Examples included predicting likely donors, estimating future donations, or using existing segmentations to anticipate behaviour in new audiences. They also emphasised the importance of practical usability, with P3 noting that "*we actually have to use what we're given.*"

Participants imagined tasks such as requesting lists of audiences who meet specific criteria. P6 gave an example of asking a system to "*identify people who booked five times a year [...] spent over this amount of money [...] give us a list of all these people so we can try and see if we can convert them to donors.*"

4.5 Healthy Mistrust of AI

Participants expressed scepticism about AI systems, especially concerning data quality and the interpretability of outputs. They emphasised the need to verify AI-generated results, understand their sources, and challenge outputs when necessary. As P6 asked, "*what if we give it bad data?*" Similarly, P3 noted the need for transparency: "*give us the sources. [...] we check each other's work before it's sent out. Who's checking the AI?*"

5 Discussion

This section brings together the findings from Section 4 to propose guiding principles for the design of AI systems that support theatre marketing and communications. The principles translate the values expressed by participants into actionable design considerations, illustrating how AI can enhance existing organisational practices while respecting the constraints and priorities of cultural institutions. To demonstrate how these principles may be put into practice, two hypothetical system concepts are presented. These examples illustrate possible pathways for implementing AI in ways that align with participants' needs, without assuming that such solutions must follow a fixed technical architecture.

5.1 Guiding Principles

The guiding principles presented in this section are grounded in the themes and subthemes identified through the Reflexive Thematic Analysis. They reflect the interplay between practical constraints, organisational values, and the types of support theatre staff considered meaningful. Together, they offer a structured foundation for future AI design in cultural organisations.

5.1.1 *Design for Workload Relief.*

AI systems should relieve workload pressures rather than introduce additional burdens. They must simplify repetitive manual tasks, streamline workflows, and fit seamlessly into existing practices. Ease of use is essential, with minimal user-training requirements and clear time-saving benefits.

To achieve this, AI systems should be designed to integrate seamlessly into existing workflows and reduce time-consuming tasks. As described in the themes, particularly within Workload Constraints Limit Flexibility, it is important to ensure that AI solutions can fit to existing workflows. Tools should minimise cognitive and temporal load, offering outputs that are quick to interpret and apply. Interpretable outputs can also aid with the concerns discussed in the themes of Healthy Mistrust of AI and Improving Data Clarity.

5.1.2 *Support Accuracy and Creativity.*

AI tools must ensure factual accuracy in details such as timings, pricing, and show information. At the same time, they should preserve space for creative expression and storytelling, which are essential to authentic audience engagement. As outlined in the Challenges in Creating Content theme, the goal is to augment, rather than replace, the human element in communications.

5.1.3 *Enable Adaptive Granularity of Audience Insights.*

Effective systems should provide insights at multiple levels of granularity. This includes aggregate audience segmentation for strategic planning, as well as personal insights for cultivating relationships with donors and frequent attendees. The ability to flexibly move between overview and detail is crucial.

AI systems must maintain a consistently accurate depiction of the data. As discussed in themes such as Granularity of Audience Understanding and Improving Data Clarity, data analysis must be consistently updated and be done to varying levels of granularity to remain accurate and impactful. Changes in audience behaviours over time and even external constraints can alter the way data may be interpreted, therefore a flexible system that is able to update is essential to provide granular audience insights.

5.1.4 *Respect Diversity in Audience Interaction.*

AI must help optimise communication schedules and channels to reflect audience preferences. It should also support targeted messaging without excluding groups or over-saturating audiences with repetitive information, as outlined by the Targeting Without Excluding theme. Inclusivity should remain a core design consideration.

Generally, systems should provide meaningful insights into demographic, behavioural, and personal preferences, while remaining sensitive to privacy and ethical considerations. AI should enhance the theatre's capacity to understand audiences, both in aggregate and at the individual level as demonstrated in themes such as Identifying Audience Trends and Understanding Individuals Personally.

5.1.5 *Prioritise Data Clarity and Integration.*

Given the fragmented and inconsistent nature of current data practices, AI systems should place data clarity at the centre of design. As presented in the Improving Data Clarity theme, systems must integrate across platforms, flag missing or inconsistent data, and help challenge human assumptions that may introduce bias. This would position AI as a tool for curating and validating information.

Additionally, as exemplified in the themes of Targeting Without Excluding, Ensuring Donor Benefits are Meaningful and Opportunities for Data-Driven Insights, AI systems may not need to be complex in nature. This includes both in relation to the outputs produced, but also in relation to the architecture. AI systems may be able to leverage outputs directly from data exploration or visualisation techniques to provides users with the insights needed. It is important to reflect on the intended use-case to appropriately design for systems that are more able to provide clarity to the data.

5.1.6 *Deliver Actionable Insights.*

Where possible, AI outputs should be directly actionable and aligned with day-to-day workflows. As outlined in themes such as Effective Audience Interaction and Opportunities for Data-Driven Insights, rather than abstract predictions, systems should generate clear and usable outputs, such as donor conversion lists or audience dashboards. Insights must be interpretable and easily exportable for practical use.

5.1.7 *Embed Transparency and Accountability.*

Trust in AI requires transparency in recommendations and accountability for errors. The Healthy Mistrust of AI theme outlines how systems should provide clear sources, justifications, and explanations for outputs. Workflows should ensure human oversight remains central, enabling sta

Furthermore, AI must support personalisation in ways that enhance relationships without infringing on privacy or trust. Designers should prioritise informed consent and ethical data use, ensuring that donor and audience information is handled responsibly. The aim is to enrich personalisation while safeguarding organisational and audience values.

5.2 Proposed AI Implementations

The following system concepts illustrate how the guiding principles can be operationalised. They are intentionally high-level and adaptable, reflecting the varied capabilities and constraints of cultural organisations. Both examples show how technical design decisions can be shaped by the values identified in the workshop.

5.2.1 Audience Segmentation System.

In accordance to the themes identified in this study and the literature at large [7, 8, 29, 46], audience segmentation is an important use-case for theatre professionals. Specifically, this refers to identifying different groups of audiences within the data. Within machine learning, this typically refers to clustering techniques [3, 16, 25, 37]. Recent advances in transformer architectures have also seen an increase in transformer approaches to clustering [10, 26, 48]. Furthermore, as discussed in Section 2, Chen and Liu [8] has also seen successful data segmentation through the use of classifier models, in particular XGBoost. However, as the author outlined, the output of such models may not be interpretable, which forms a contrast to the guiding principles outlined in this paper. Similarly, transformer architectures are typically resource intense and has higher requirements of training data, which may not always be available for smaller venues.

Thus, to design a system that segments audience data, this paper proposes a clustering approach, specifically following an archetypal based strategy such as K-Maxoids [3]. This is because archetypal based approaches are known for offering greater interpretability, making it a powerful approach for quickly conveying the meaning of the clusters [9]. Online approaches of K-Maxoids [39] can also serve to offer greater granularity and adaptability to a change in audience behaviour. Importantly, as a centroid based strategy, in which the centroids are the archetypes, interpretability of the clusters can be gained by studying the archetypes directly, as a consequence of the algorithm design [3].

In regards to the user needs, multiple levels of clustering would be required. This section proposes some recommendations of different ways of how the data may be clustered. Note that in doing so, this section also provides some recommendation for data features. However, other features may be required to enable effective data segmentation. An exploration of the data is recommended during implementation of such solutions.

Purchasing Behaviour: clustering that focuses on the way in which audiences purchases tickets to shows and events. This should focus on the quantity and frequency of purchases. As a result, it is recommended that input vectors represent individual audience members, with features such as *[mean number of tickets acquired in a single purchase, total number of purchases, mean time elapsed between purchases, mean yearly tickets purchased]*.

Content Behaviour: clustering that focus on the types of content purchased by audiences. This would be focused on the genres and sub-genres of the events being purchased. Similarly, this cluster could also provide insights over the type of events (such as touring shows vs. locally produced content). To avoid overfitting to those who purchase tickets often, it is recommended to limit the analysis to a window based on number of purchases. The recommended features in this case would include *[(per genre) count of main genre purchases of the previous 5 purchases, (per sub-genre) count of sub-genre purchases of the previous 5 purchases, (per existing event type) count of event type purchases of the previous 5 purchases]*.

Content Behaviour (feedback): in some cases, audiences may also provide feedback post-attendance. As a result, the insights provided by the *Content Behaviour* may also be enhanced by analysing the (quantitive) feedback into the clusters. For example, if an audience is asked to rate their enjoyment in a 5-point Likert scale, the average responses by that individual per genre, sub-genre and event type could reveal valuable insights about audiences. This cluster type has been separated from the standard *Content Behaviour* cluster type as not all audiences are likely to feel out feedback responses. This separation offers greater flexibility.

Location: cluster based on the location of the audience member. This could be done based on either distance to the theatre directly, or simply based on where the individuals are located. In this example, address is used to create the input vector of *[latitude, longitude]*. Importantly, this clustering layer may be redundant, as addresses are likely to result in clusters based on existing cities and towns. An exploration of the results, ideally involving qualitatively or otherwise evaluating theatre professionals reception to the results, should be made to investigate.

In addition to the cluster types described, theatre professionals may also benefit from the ability to navigate through their data. Therefore, the proposed system would consist of collating their data into a single database, which would perform the clustering of audiences during the input period. While most optimal, seamlessly integrating the system into the theatre's workflows to replace existing data collection and storage solutions may be costly. One alternative that still offers seamless integration is to support custom parsers for data input that should be custom built to work alongside existing tools and protocols for data collections. In this case, data exports from existing tools (such as a password protected Excel format), would be created by users to then be read by the proposed AI system. Clustering and navigation should be done locally and no long term storage should be provided in this case, to avoid risks of data duplication and security vulnerabilities. Data should only be processed and stored using volatile storage using the input files provided by the users. Users should be given the option to automatically delete input files to also reduce security threats.

In order to navigate the data, cluster types should be treated as additional features for individual users. Within the database, this would be represented as a table, which links audience ID to its respective cluster for each of the cluster types discussed.

Users should then be given the ability to query the data. To enable ease of usability, this is likely best achieved through intuitive design principles relying on drop down menus and input fields as depicted in Figure 1. While not the primary way that users are expected to interact with the system, it may be beneficial to support custom SQL queries, to enable more experienced users to take advantage of the system to its maximum.

Within the wireframe in Figure 1, the "Advance Search (SQL)" button would allow power users to enter a SQL query manually. The "View Archetypes" button would open a report that describes all of the archetypes for each of the cluster types proposed. This description would include the archetype name, the de-normalised

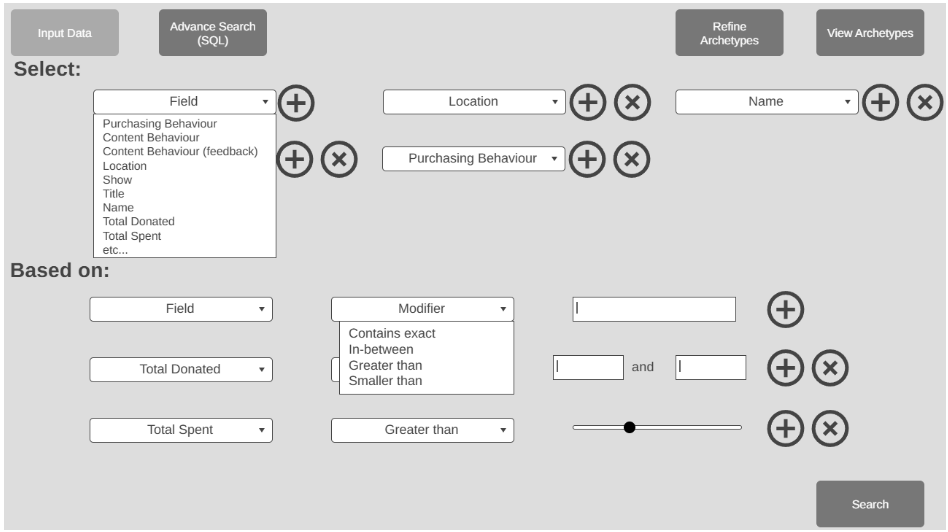

Figure 1: A wireframe depicting how an AI system for audience segmentation could be presented to users for utilisation

values for the centroids and a description for the archetype. These would be refined through the use of the "Refine Archetypes" button, which would use the Online K-Maxoid approach [39] to use the dataset provided to further refine the data segmentation. During this process, users would need the ability to see changes in the values for the centroids, giving them the option to alter the archetype's name and description.

To design this system in relation to the guiding principles, the following is considered:

- *Design for Workload Relief*: Audience segmentation is already a current need of theatre professionals, which is currently being done manually through their understanding of the data. This is a time consuming process, which is also vulnerable to inject human bias in the output. Producing an automated system would expedite the process of completing existing tasks without introducing additional burden on professionals.
- *Support Accuracy and Creativity*: With the way in which the system is designed, theatre professionals can choose to query their own lists based on their preferences and requirements. This allows them to use creative ways of investigating the data, while providing them with accurate data-driven insights.
- *Enable Adaptive Granularity of Audience Insights*: By providing users with multiple levels of clustering, theatre professionals can take advantage of granular insights which reflect the multifaceted nature of audience behaviours and preferences.
- *Respect Diversity in Audience Interaction*: By enabling continued refinement of the data segmentation, the AI model can be sensitive to audience diversity and changes in they way they interact with the theatre. Similarly, the explorative nature of the solution also allows for complex identification of diverse groups, not only based on the AI insights, but also by combining those with specific data features that are available in the dataset.
- *Prioritise Data Clarity and Integration*: By providing custom built parsers that can directly integrate with existing workflows and tools, the proposed AI system would provide more clarity to their existing data, in addition to AI driven insights. This promotes better experiences for analysing data in a centralised solution.
- *Deliver Actionable Insights*: By enabling flexibility of data explorations, users may more effectively target audiences. As the solution would generate a list of customers, including any desired fields, practitioners may use the insights to communicate with audiences directly, offering clear and actionable insights from the data.
- *Embed Transparency and Accountability*: By combining the interpretability of archetypes with a detailed report depicting and describing the archetypes, the system provides a clear and transparent depiction of how the insights are generated. Furthermore, the by maintaining users in the loop, with the refinement of archetypes over time, the system can remain accountable as clear descriptions as well as use-case strategies are developed over time.

5.2.2 Donor Identification System.

Donor identification is another example use-case that participants have discussed during the co-design workshop. This system could contain two tools. Firstly, direct donor identification, in which AI is used to predict (based on the current audiences who do not already donate), who are the most likely individuals to become donors. Secondly, this system could provide a break-down of the likely amount of donations that the theatre is to receive based on the shows.

To achieve donor identification, archetypal clustering may not be sufficient by itself. Instead, this solution should pair a predictor model, such as a predictive neural network (NN), with archetypes. In doing so, the NN can identify which audience members are likely to become donors, while donor archetypes may give insights over what type of donor they may become.

The NN architecture in this case could attempt to predict the donation amount of each individual, including 0 (no donation). To predict this, data about current donors would need to be used as training data. This could include input features such as *[previous donation amount, average donation amount, mean time elapsed between donations]* combined with either the archetypes described in Section 5.2.1 or the same features discussed. The output of such model could be the predicted donation amount for each individual. In relation to the training process, representative data for each donor at the time of each donation would be used as the training input dataset (X), with the donation amount as the label (Y).

Importantly, such model is not transparent in relation to its outputs. To aid in transparency, and to provide actionable insights, archetypes for the donor types can be generated using the same methodology described in Section 5.2.1. This ensures the users of the system can maintain trust and agency over the outputs. To train such archetypes, K-Maxoids would be used with input features such as *[total donation amount, average donation amount, mean time elapsed between donations, most recent donation amount, number of donations]*.

To present the output of these models, a system that generates reports should be developed. This could be integrated into the system depicted in Figure 1, with the addition of a button to view the report. This report is depicted in Figure 2.

Within the donor identification system, the system could use knowledge of monthly donors (who are part of monthly programs already) in addition to the predicted amounts expected by the NN to approximate the expected amount of donations to be received on each month. The "See breakdown" button would generate a report identifying the current donors that are already signed up to donate for the month and how much this total is. The report would then also highlight the predicted amount. In a concrete implementation, the total amount depicted in Figure 2, could be colour coded (or though other UI design) to quickly and intuitively differentiate between secure amount (from those who are already in monthly donation schemes) and predicted amount (the result from the NN).

As with the archetypes described in Section 5.2.1, the donor archetypes can be explained and further refined within the system depicted in Figure 2. The "See Donor Archetypes" button would provide users with a report depicting all of the archetypes, including names, descriptions and centroid values. The "Refine Donor Archetypes button would then use the online K-Maxoids approach [39] to refine the centroids, a process which would allow users to compare, rename and alter the descriptions for each of the new archetypes if required.

The "Current Donor" and "Predicted Donor" buttons would serve as toggles, changing which group is being presented in the report detail report. This report would provide a range of meaningful information about the donor as an individual, such as contact information, donor archetype as classified with the clustering model and previous donation information. Additionally, if integrated alongside the system described in Section 5.2.1, the various archetypes describes would also be displayed in the detailed report.

Lastly, the "Export List" button would produce a external file with the donor information. This would likely best achieved through a password encrypted file to ensure data security. The exported list would need to be generated following custom rules to allow for best integration to any existing systems used by theatres. This would ensure outputs and insights generated by the tools can be leverage by theatre professionals existing workflows.

To design this system in relation to the guiding principles, the following is considered:

- *Design for Workload Relief*: By producing reports through simple interactions, this system would allow for easier donor targetting without adding any time consuming steps. As donor identification and communication is already a task that theatre marketing and communications professionals undertake, this tools is designed to ease the pain points associated with completing this workload. This focuses on expediting existing workflows for donor identification, content creation and communication targetting.
- *Support Accuracy and Creativity*: To enable deeper understanding and support creative use of data, the pairing of the clustering and predictive NN allows for insights to remain accurate, data-driven but also offer professionals the freedom to utilise the insights in their preferred way. Particularly, the use of donor archetypes can support communication to be targetted in more unique ways, ensuring creative expression is supported.
- *Enable Adaptive Granularity of Audience Insights*: By providing a breakdown discussion of predicted and expected values, this system can provide greater granularity. Additionally, not only are the predicted donors identified, but the further characterisation of donor archetypes and the predicted donation amount provide greater detail in a granular manner.
- *Respect Diversity in Audience Interaction*: By enabling continued refinement of the data segmentation, the AI model can be sensitive to audience diversity and changes in they way they interact with the theatre.
- *Prioritise Data Clarity and Integration*: By providing custom built parsers that can directly integrate with existing workflows and tools, the proposed AI system would provide more clarity to their existing data, in addition to AI driven insights. This promotes better experiences for analysing data in a centralised solution.
- *Deliver Actionable Insights*: The system produces a list of individuals to be targetted in communication pieces. This provides a clear action to be taken by theatre professionals. additionally, the various types of donors, both existing and predicted, which are described by the archetypes can give insights into how to best to communicate with them in regards to content generation.
- *Embed Transparency and Accountability*: As the proposed systems includes a black-box system, transparency is mitigated by the way in which insights are reported and through maintaining users in the loop, such as through the refinement of the archetypes. It is important to keep the user informed of the non-transparent nature of the predicted donation amount, which places a higher importance in the generation of the detailed report which aids in maintaining transparency of the overall system.

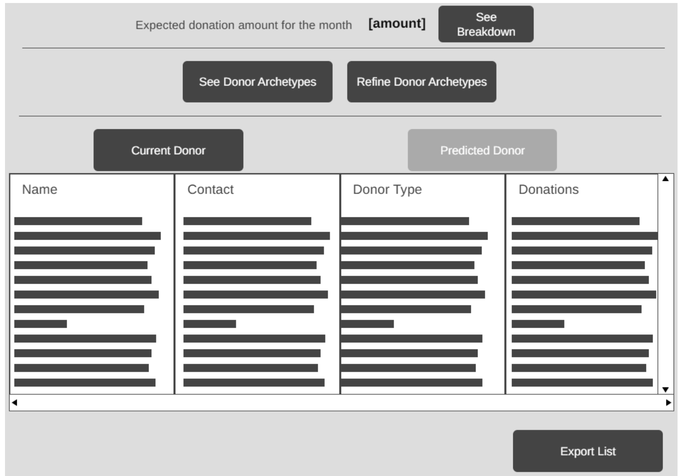

Figure 2: A wireframe depicting how an AI system for donor identification could be presented to users for utilisation

5.3 Situating AI Principles

As the proposed AI implementations depict illustrative examples of how the guiding principles may be operationalised, it is important to situate these principles within the broader technological ecosystem in which theatre organisations currently operate. In particular, existing CRM systems represent a key point of comparison, given their widespread adoption across theatre marketing and development teams.

Table 1 outlines how the guiding principles identified in this study align with, and diverge from, the capabilities of commonly used CRM platforms. Rather than evaluating specific products, the table synthesises limitations described by participants and observed across the findings, illustrating where current systems fall short of supporting day-to-day communication and interpretive work. This comparison highlights the gap addressed by the proposed principles and clarifies the contribution of this paper within existing organisational infrastructures.

Importantly, this comparison is not intended as a critique of CRM systems, but as an illustration of the gap between their original design intent and the values articulated by theatre communications professionals in this study. By situating the guiding principles alongside established tools, this paper demonstrates how its contributions extend existing capabilities rather than duplicate them, and why new AI approaches are needed to better support communicative practice in theatre contexts.

6 Limitation & Future Work

This study offers detailed insight into the practices and values of a single theatre's marketing and communication team. While this provides strong contextual grounding, values and workflows may vary across other cultural organisations. Different types of arts and cultural organisations beyond theatres and concert halls, such as museums or charities remain largely underexplored and, while some generalisability is expected, continued research in the domain is necessary to better understand their specific needs. Additionally, within the domain of theatres and concert halls, organisations that primarily cater for a different populations, be it geographical or social-economical, may also have varying needs. Future research should therefore explore how the guiding principles translate across institutions of different sizes, regions, and organisational structures.

The AI system concepts presented here illustrate how the principles may be operationalised, but they remain hypothetical. Developing and deploying such systems in real settings would reveal further constraints and opportunities, particularly regarding data availability, ethical considerations, and long-term integration into organisational practices.

Overall, the guiding principles derived in this study offer a foundation for designing AI that meaningfully supports cultural sector work, emphasising usability, transparency, and alignment with organisational values.

7 Conclusion

This paper examined how AI can be designed to support theatre marketing and communications, a domain that has received limited attention despite its importance to sustaining cultural institutions. Through a co-design workshop with the full marketing and communications team at Perth Theatre and Concert Hall, the study identified values, needs, and concerns that shape professionals' day-to-day work. Using Reflexive Thematic Analysis, these insights were synthesised into a set of themes that capture how staff navigate workload pressures, data challenges, audience engagement, and emerging expectations around AI.

Building on these findings, the paper introduced a set of guiding principles for the design of AI systems that align with theatre professionals' values and practices. These principles emphasise workload relief, accuracy and creativity, adaptive granularity, inclusivity in audience engagement, data clarity, actionable insight, and transparency. They offer a structured foundation for developing AI that integrates meaningfully into organisational workflows rather than disrupting them.

To illustrate how these principles may be operationalised, the paper presented two hypothetical system concepts: an audience

Table 1: Alignment between existing CRM platforms and design principles derived from the workshops.

Design Principle	Current CRM Capabilities	Observed Limitations in Practice
Workflow relief and integration	Automates ticketing, donations, mailing lists, and financial reporting.	Limited support for the creative and interpretive work of crafting communications, requiring manual translation of analytics into action across multiple tools.
Adaptive audience understanding	Provides segmentation based on attendance history, demographics, and donation behaviour.	Supports primarily aggregate views, offering limited flexibility for incorporating contextual, relational, or situational knowledge valued by staff.
Effective and inclusive communication	Enables bulk email campaigns and basic targeting.	Offers limited assistance with balancing personalisation, inclusivity, and communication fatigue, leaving ethical judgement entirely to staff.
Data clarity and transparency	Supplies dashboards and performance metrics.	Data provenance, completeness, and bias are often opaque, particularly when data are fragmented across platforms.
Actionable insights	Increasingly incorporates predictive analytics for revenue optimisation.	Outputs prioritise retrospective reporting or revenue-focused prediction rather than supporting day-to-day communication decisions.
Transparency and accountability	Provides stable, auditable transaction records.	AI-driven features offer limited explanation or contestability, reducing staff confidence in applying outputs to communicative practice.

segmentation tool based on archetypal clustering, and a donor identification tool combining predictive modelling with interpretable archetypes. These examples demonstrate how AI could support existing tasks while remaining sensitive to practical constraints, ethical considerations, and the need for interpretability.

Taken together, this work contributes an empirically grounded understanding of theatre professionals' needs and provides a principled basis for designing AI that supports cultural sector practices. Future research can expand this work by exploring diverse organisational contexts and by developing and evaluating AI systems informed by the principles proposed here.

Acknowledgments

This research was part-funded by the UKRI Arts and Humanities Research Council funded CoSTAR Live Lab project (AH/Y001079/1) and by the UKRI Engineering and Physical Sciences Research Council Centre for Doctoral Training in Intelligent Games & Games Intelligence (IGGI) (EP/S022325/1).

This project was supported by Perth Theatre and Concert Hall. The authors would like to extend a special thanks to Christopher Glasgow for facilitating the workshop and for all of the valuable insights they brought up. Similarly, the authors would also like to thank the other members of staff from Perth Theatre and Concert Hall for their participation in the workshop, their expertise and experience in the field allowed for an in-depth exploration without which this paper would not have been possible.

References

[1] Kenneth C. Arnold, Krysta Chauncey, and Krzysztof Z. Gajos. 2020. Predictive text encourages predictable writing. In *Proceedings of the 25th International Conference on Intelligent User Interfaces* (Cagliari, Italy) *(IUI '20)*. Association for Computing Machinery, New York, NY, USA, 128–138. doi:10.1145/3377325.3377523

[2] Federico Avanzini, Adriano Baratè, Goffredo Haus, Luca A. Ludovico, and Stavros Ntalampiras. 2020. Preservation and Promotion of Opera Cultural Heritage: The Experience of La Scala Theatre. In *Culture and Computing*, Matthias Rauterberg (Ed.). Springer International Publishing, Cham, 325–337.

[3] Christian Bauckhage and Rafet Sifa. 2015. k-Maxoids Clustering.. In *n Proceedings of the LWA 2015 Workshops: KDML*. 133–144.

[4] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative research in sport, exercise and health* 11, 4 (2019), 589–597.

[5] Virginia Braun and Victoria Clarke. 2021. *Thematic analysis: A practical guide.* SAGE publications Ltd.

[6] Niccole Carner. 2018. Building national identity and cultural confidence in the National Theatre of Scotland's theatre for young audiences. *Youth Theatre Journal* 32, 2 (2018), 138–146.

[7] Qingjiao Chen, Yaxian Wang, and Zhi Tang. 2025. Impact Mechanism of AI-Driven Content Marketing on Purchase Intentions: The Moderating Role of Education. In *Proceedings of the 2nd Guangdong-Hong Kong-Macao Greater Bay Area International Conference on Digital Economy and Artificial Intelligence.* 1226–1231.

[8] Yuezhang Chen and Jiacheng Liu. 2025. AI-Driven Marketing Strategy Optimization in the Digital Economy: A Machine Learning Approach. In *Proceedings of the 2025 International Conference on Digital Economy and Intelligent Computing.* 239–242.

[9] Alan Pedrassoli Chitayat. 2025. AI vs. the Algorithm: Measuring Success on Twitch. In *2025 IEEE Conference on Games (CoG)*. IEEE, 1–8.

[10] Kenneth L Clarkson, Lior Horesh, Takuya Ito, Charlotte Park, and Parikshit Ram. 2025. Finding Clustering Algorithms in the Transformer Architecture. *arXiv*

preprint arXiv:2506.19125 (2025).

[11] Ankolika De and Zhicong Lu. 2024. # PoetsOfInstagram: navigating the practices and challenges of novice poets on Instagram. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–16.

[12] Erika Fischer-Lichte. 1996. IN CONTEMPORARY THEATRE. *The intercultural performance reader* (1996), 27.

[13] Batya Friedman. 1996. Value-sensitive design. *interactions* 3, 6 (1996), 16–23.

[14] Batya Friedman, Peter Kahn, and Alan Borning. 2002. Value sensitive design: Theory and methods. *University of Washington technical report* 2, 8 (2002), 1–8.

[15] Brett A Halperin, Diana Flores Ruíz, and Daniela K Rosner. 2025. Underground AI? Critical approaches to generative cinema through amateur filmmaking. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–18.

[16] Greg Hamerly and Charles Elkan. 2003. Learning the k in k-means. *Advances in neural information processing systems* 16 (2003).

[17] Hugo Hammond, Michael Armstrong, Graham A Thomas, and Iain D Gilchrist. 2023. Audience immersion: validating attentional and physiological measures against self-report. *Cognitive Research: Principles and Implications* 8, 1 (2023), 22.

[18] Tharon Howard. 2014. Journey mapping: A brief overview. *Communication Design Quarterly Review* 2, 3 (2014), 10–13.

[19] Maurits Kaptein and Petri Parvinen. 2015. Advancing e-commerce personalization: Process framework and case study. *International Journal of Electronic Commerce* 19, 3 (2015), 7–33.

[20] Turkka Keinonen. 2008. User-centered design and fundamental need. In *Proceedings of the 5th Nordic conference on Human-computer interaction: building bridges*. 211–219.

[21] Ahmed Khan. 2024. Performing History: The Role of Theatre in Preserving Cultural Heritage. *Al-Zumar* 2, 01 (2024), 1–8.

[22] Xuye Liu, Annie Sun, Pengcheng An, Tengfei Ma, and Jian Zhao. 2025. Influencer: Empowering Everyday Users in Creating Promotional Posts via AI-infused Exploration and Customization. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–19.

[23] Piotr Mirowski, Kory W Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-writing screenplays and theatre scripts with language models: Evaluation by industry professionals. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–34.

[24] Meena Devii Muralikumar and David W McDonald. 2024. Analyzing Collaborative Challenges and Needs of UX Practitioners when Designing with AI/ML. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW2 (2024), 1–25.

[25] Fionn Murtagh and Pedro Contreras. 2012. Algorithms for hierarchical clustering: an overview. *Wiley interdisciplinary reviews: data mining and knowledge discovery* 2, 1 (2012), 86–97.

[26] Xuan-Bac Nguyen, Duc Toan Bui, Chi Nhan Duong, Tien D Bui, and Khoa Luu. 2021. Clusformer: A transformer based clustering approach to unsupervised large-scale face and visual landmark recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10847–10856.

[27] Richard A Oakes, Lisa Peschel, and Nick E Barraclough. 2024. Inter-subject correlation of audience facial expressions predicts audience engagement during theatrical performances. *IScience* 27, 6 (2024).

[28] Alan Pedrassoli Chitayat, Florian Block, James A Walker, and Anders Drachen. 2024. How Could They Win? An Exploration of Win Condition for Esports Narratives in Dota 2. *Proceedings of the ACM on Human-Computer Interaction* 8, CHI PLAY (2024), 1–22.

[29] Pythagoras Petratos and Mina Giannoula. 2024. Transformer-based AI for Sentiment Analysis in Marketing. In *Proceedings of the 2024 6th Asia Conference on Machine Learning and Computing*. 177–185.

[30] Yim Register, Lucy Qin, Amanda Baughan, and Emma S Spiro. 2023. Attached to "The algorithm": Making sense of algorithmic precarity on instagram. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–15.

[31] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. " Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–1144.

[32] Raquel Breejon Robinson, Ricardo Rheeder, Madison Klarkowski, and Regan L Mandryk. 2022. " Chat Has No Chill": A Novel Physiological Interaction For Engaging Live Streaming Audiences. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–18.

[33] Mark S Rosenbaum, Mauricio Losada Otalora, and Germán Contreras Ramírez. 2017. How to create a realistic customer journey map. *Business horizons* 60, 1 (2017), 143–150.

[34] Malak Sadek, Marios Constantinides, Daniele Quercia, and Celine Mougenot. 2024. Guidelines for integrating value sensitive design in responsible AI toolkits. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–20.

[35] Hooman Samani, Vali Lalioti, Diana Alina Serbanescu, Joana Chicau, Doros Polydorou, George Rodosthenous, Amelia Knowlson, Yorgos Bakalos, Michael Neale, and Bipin Indurkhya. 2024. Creative robotics theatre: designing creative interactions with tangible and embodied interfaces. In *Companion Publication of the 2024 ACM Designing Interactive Systems Conference*. 389–391.

[36] Lubov G Savenkova, Olga I Radomskaya, Inga V Zhgenti, et al. 2019. Theatre Art In The Information Age: Youth Audience Perception Features. *European Proceedings of Social and Behavioural Sciences* (2019).

[37] Erich Schubert, Jörg Sander, Martin Ester, Hans Peter Kriegel, and Xiaowei Xu. 2017. DBSCAN revisited, revisited: why and how you should (still) use DBSCAN. *ACM Transactions on Database Systems (TODS)* 42, 3 (2017), 1–21.

[38] Katie Shilton et al. 2018. Values and ethics in human-computer interaction. *Foundations and Trends® in Human–Computer Interaction* 12, 2 (2018), 107–171.

[39] Rafet Sifa and Christian Bauckhage. 2017. Online k-maxoids clustering. In *2017 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. IEEE, 667–675.

[40] Marc Stickdorn and Jakob Schneider. 2012. *This is service design thinking: Basics, tools, cases.* John Wiley & Sons.

[41] Milka Trajkova, Duri Long, Manoj Deshpande, Andrea Knowlton, and Brian Magerko. 2024. Exploring collaborative movement improvisation towards the design of LuminAI—a co-creative AI dance partner. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–22.

[42] Wolfgang Tschacher, Steven Greenwood, Hauke Egermann, Melanie Wald-Fuhrmann, Anna Czepiel, Martin Tröndle, and Deborah Meier. 2023. Physiological synchrony in audiences of live concerts. *Psychology of aesthetics, creativity, and the arts* 17, 2 (2023), 152.

[43] Wolfgang Tschacher, Steven Greenwood, Sekhar Ramakrishnan, Martin Tröndle, Melanie Wald-Fuhrmann, Christoph Seibert, Christian Weining, and Deborah Meier. 2023. Audience synchronies in live concerts illustrate the embodiment of music experience. *Scientific Reports* 13, 1 (2023), 14843.

[44] Robert Twomey, Ash Eliza Smith, Reid Brockmeier, and Samantha Bendix. 2025. Quantum Theater: Extending Realities for Post-AI Liveness. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Spatial Storytelling*. 1–3.

[45] Alexandra Tzanidou, Al Husein Sami Abosaleh, Stephen Lindsay, Abigail C Durrant, and Vasilis Vlachokyriakos. 2023. From Inclusive Theatre to inclusive technologies: Lessons learnt from co-designing Touch Tours with an Inclusive Theatre group. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 1367–1382.

[46] Kulnapruek Vongpanich, Kanyaphas Boonyapitaktumrong, Cherliya Veerathanusvet, and Chutisant Kerdvibulvech. 2025. The Impact of AI on Enhancing Productivity in Digital Marketing Production. In *Proceedings of the 2025 7th Asia Pacific Information Technology Conference*. 62–66.

[47] Ruyuan Wan, Lingbo Tong, Tiffany Knearem, Toby Jia-Jun Li, Ting-Hao'Kenneth' Huang, and Qunfang Wu. 2025. Hashtag re-appropriation for audience control on recommendation-driven social media Xiaohongshu (rednote). In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–25.

[48] Ningning Wang, Guobing Gan, Peng Zhang, Shuai Zhang, Junqiu Wei, Qun Liu, and Xin Jiang. 2022. ClusterFormer: Neural clustering attention for efficient and effective transformer. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 2390–2402.

[49] James G Webster. 2016. *The marketplace of attention: How audiences take shape in a digital age.* MIT press.

[50] Joseph Williams, Jon Francombe, and Damian Thomas Murphy. 2024. Camera Sourced Heart Rate Synchronicity: A Measure of Immersion in Audiovisual Experiences. *Applied Sciences* (2024).

[51] Sojeong Yun and Youn-kyung Lim. 2025. User Experience with LLM-powered Conversational Recommendation Systems: A Case of Music Recommendation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–15.

[52] Jiazhi Zhou, Rebecca Weber, Elliott Wen, and Danielle Lottridge. 2025. Real-Time Full-body Interaction with AI Dance Models: Responsiveness to Contemporary Dance. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*. 1177–1187.

A Co-design Workshop Artifacts

Note that after the conclusion of the workshop artifacts were labelled by the research team with (A) and (B), relating to which group the artifacts were from. This was only done after, purely to aid in record keeping.

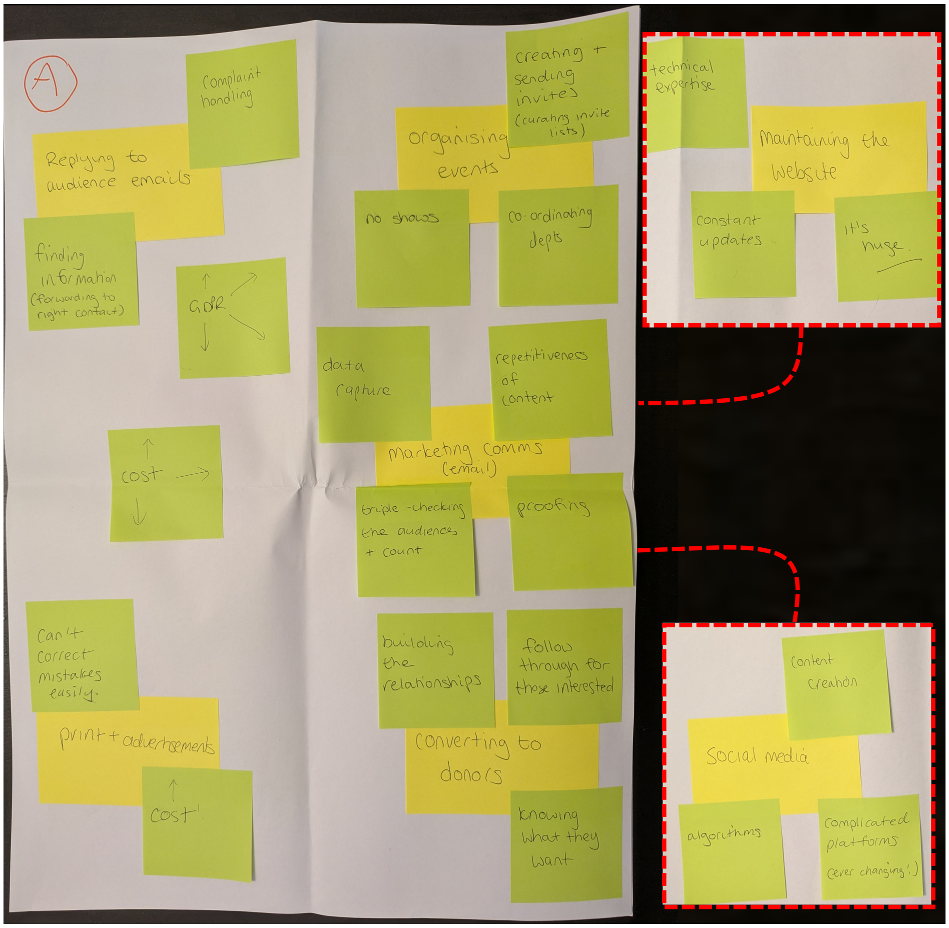

Figure 3: The artifacts generated by Group A (P5-P8) during Task 1

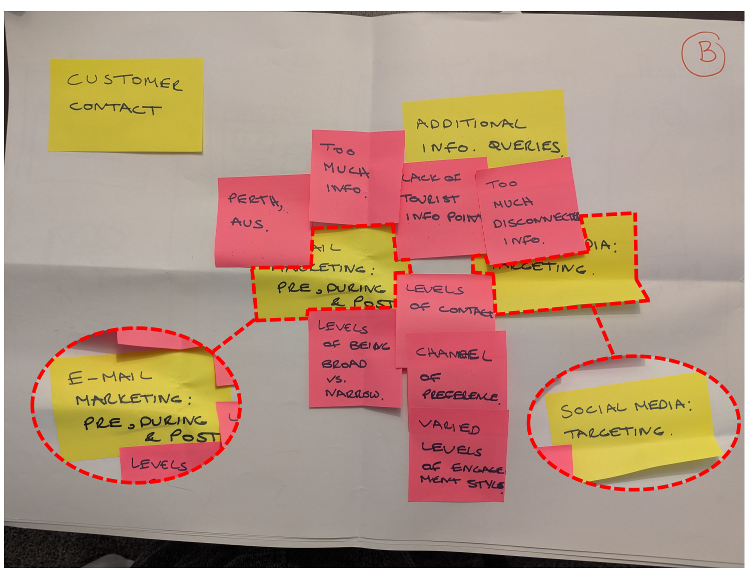

Figure 4: The artifacts generated by Group B (P1-P4) during Task 1

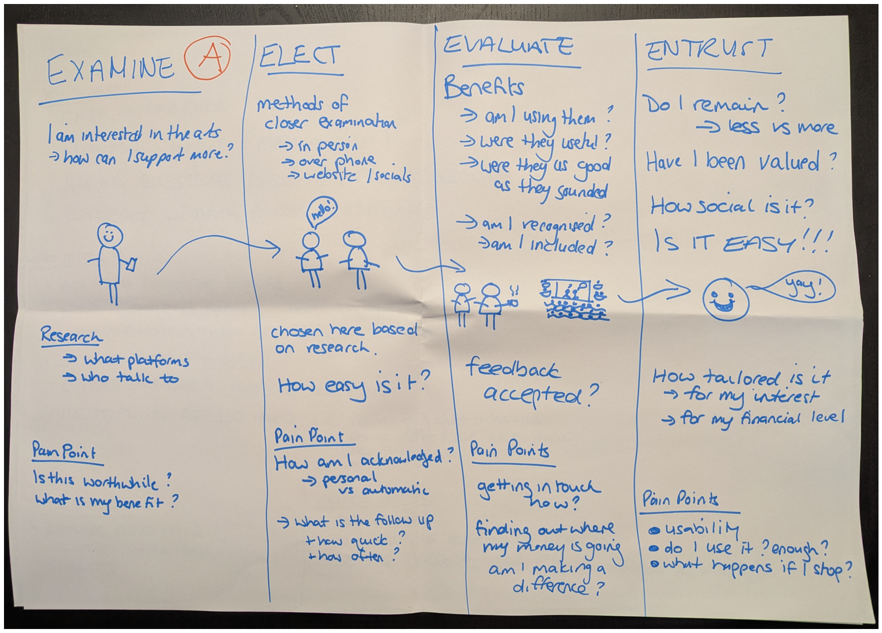

Figure 5: The user journey generated by Group A (P5-P8) during Task 2

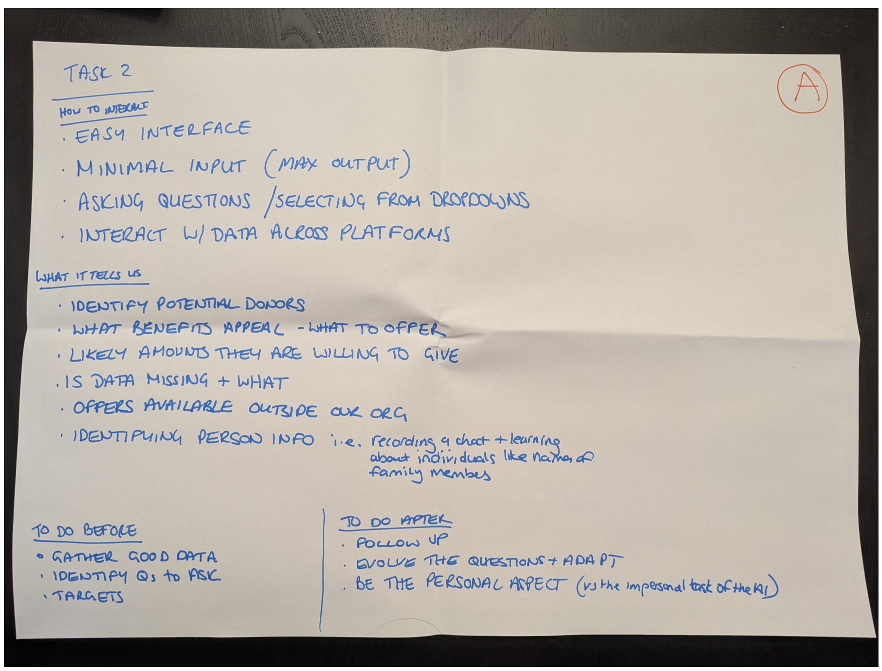

Figure 6: Supplementary notes generated by Group A (P5-P8) during Task 2