Author/s:
Balkin, S;Walton, RE

Title:
Can AI read the room?: Attuning machines to the affective atmospheres of stand-up comedy performance

Date:
2025-01-01

Citation:
Balkin, S. & Walton, R. E. (2025). Can AI read the room?: Attuning machines to the affective atmospheres of stand-up comedy performance. Comedy Studies, 16 (2), pp.314-332. https://doi.org/10.1080/2040610X.2025.2544414.

Persistent Link:
https://hdl.handle.net/11343/357698

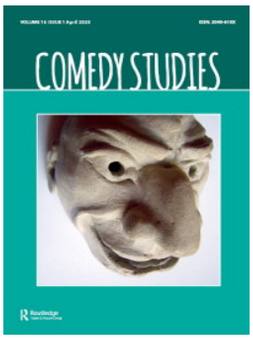

Comedy Studies

ISSN: 2040-610X (Print) 2040-6118 (Online) Journal homepage: www.tandfonline.com/journals/rcos20

Can AI read the room?: Attuning machines to the affective atmospheres of stand-up comedy performance

Sarah Balkin & Robert Ellis Walton

To cite this article: Sarah Balkin & Robert Ellis Walton (12 Aug 2025): Can AI read the room?: Attuning machines to the affective atmospheres of stand-up comedy performance, Comedy Studies, DOI: 10.1080/2040610X.2025.2544414

To link to this article: https://doi.org/10.1080/2040610X.2025.2544414

COMEDY STUDIES
https://doi.org/10.1080/2040610X.2025.2544414

Routledge
Taylor & Francis Group

Can AI read the room?: Attuning machines to the affective atmospheres of stand-up comedy performance

Sarah Balkin[a] and Robert Ellis Walton[b]

[a]English and Theatre Studies, University of Melbourne, Melbourne, Australia; [b]Faculty of Fine Arts and Music, University of Melbourne, Melbourne, Australia

ABSTRACT

This article explores the potential and limitations of artificial intelligence (AI) for 'reading the room' in stand-up comedy, a genre traditionally reliant on very human skills: comedians must acknowledge the performance situation and respond to affective atmospheres in the moment to maintain their license to speak. As AI blurs the boundaries between human and machine capabilities, this research investigates how it might analyse and engage with both live and digital comedy audiences. Drawing on Henri Bergson's concept of laughter as 'something mechanical encrusted on something living,' the article employs a multifaceted approach, analysing AI's current capabilities in interpreting recorded performances, speculating on future applications using sensors, and considering ethical implications of AI usage. While AI shows promise in analysing recorded performances and audience data, it faces significant challenges in replicating the nuanced, real-time skills of human comedians. The article considers the potential for AI to supplement or surpass human capabilities in certain aspects of audience analysis but highlights persistent limitations in developing a believable comedic persona. This study understands AI's role in stand-up comedy as augmentative rather than replacive, suggesting new directions for AI comedy research, including the exploration of 'more-than-human' forms of comedy and the adaptation of live performance training methods for AI.

KEYWORDS

Audience studies; stand-up comedy; artificial intelligence (AI)

The central image of French philosopher Henri Bergson's essay 'On Laughter' (1900) is his formulation of the comic as 'something mechanical encrusted on something living' (2003, 25). For Bergson the mechanical aligns with rigidity of body, mind, or character (2003, 18) whereas vitality aligns with an elasticity that enables people to adapt to the situation at hand (2003, 17). Laughter is the social corrective to the rigidities that limit human adaptability (Bergson 2003, 18). In other words, a person is funny when they act 'more like a machine' and comedy's role 'is to draw the audience's attention to this image' to bring them back in touch with humanity's vitality and adaptability (Milner Davis 2014, 79). Further, 'the comic does not exist outside the pale of what is strictly HUMAN' (Bergson 2003, 11, capitalisation in original).

CONTACT Sarah Balkin sarah.balkin@unimelb.edu.au Senior Lecturer in English and Theatre Studies, University of Melbourne, Melbourne, Australia.

Although we may laugh at an animal or an object, Bergson argues, what we are really laughing at is 'some human attitude or expression' in the animal or 'the human caprice whose mould' the object resembles (2003, 11).[1]

While Bergson's theory of laughter continues to influence comedy scholarship, expanded definitions of the human and our relationships to technology complicate his formulation of the comic as something mechanical encrusted on something living. For example, building on the insights of object-oriented ontology (OOO), Sarah Lucie rejects binary oppositions between the material human body and the seemingly immaterial datasphere, even as she argues that the body resists complete absorption in the technosphere. In twenty-first-century media, Lucie notes, 'a human's qualities are translated into data that goes on to interact in algorithmic technology and machine-learning systems in particular', creating 'an expanded human object that is distributed in time and space' (2021, 320). Put simply, the emergence of artificial intelligence (AI) is eroding (but not obliterating) 'the boundaries between human and computer' (Metzger 2021, xv) and increasingly, digital data is mechanically encrusted on contemporary life to describe our habits: whether we like them or not, these data trace our rigidities and machine-like peculiarities that make our behaviour predictable to algorithms. This article examines the implications of this erosion of boundaries and encrustation of data for stand-up comedy, a genre that traditionally requires very human skills: comedians must acknowledge the performance situation and respond to affective atmospheres in the moment to maintain their license to speak. Where existing research on AI's applications for comedy tends to focus on its capabilities and limitations as a tool for writing it, our focus is on the still-distant frontier of reading a comic audience.[2]

The skillset of responding to an audience in real time is colloquially called 'reading the room'. The room is both a physical space where the performance takes place and a metonymy for the audience and its reactions (Balkin and Mierowsky 2024, 39). Reading the room is the comedian's job, but it is interactive rather than one-sided, especially in a more intimate venue (Mintz 1985, 78). The live venue in which a performance takes place influences the audience's attention, emotions, and behavior (Quirk 2015, 65-77). With the increased production and circulation of stand-up for and in digitised forums such as streaming services and social media platforms, comedy scholars have begun to reconceive who and what is in the expanded 'room' and what that means for the relationship between performers and audiences. Many high-profile comedians now speak to more than one audience in their shows, addressing the people congregated in the live room; an expanded, dispersed streaming platform audience; and social media users and popular media commentators who often re-circulate text, images, and video from digitised comedy shows in ways that strip or alter the performance context (Balkin 2023, 158; Balkin and Mierowsky 2024, 36-47; Goltz 2017, 94). This article builds on key ideas from contemporary audience studies—namely, that the audience is active as well as receptive; that the audience's conventional behavior and affective responses are historically situated and embodied; and that the audience co-creates the performance event and its meaning—to imagine how AI might read this expanded comedy audience.[3] As digital technologies become more pervasive, humans 'inhabit environments increasingly structured through algorithms' (Metzger 2021, xvi). These data networks 'integrate AI and people' in ways that distribute humans across media platforms, extending Marshall McLuhan's aspiration

for technology to 'extend our central nervous system itself in a global embrace' (1964, 16) and eroding 'the boundaries between human and computer' (Metzger 2021, xv). Sean Metzger calls these new relations between people and technologies 'digital cohabitation', where machines learn from users and users learn from machines, each developing what [Wendy Hui Kyong] Chun calls 'habits' that facilitate person and computer's ostensible mutual understanding' (2021, xv). If the room has expanded into digital spaces, this article asks, can AI read it, and what might that mean?

The use of AI to analyse and create comedy raises ethical concerns that include: the provenance of training data that built the weighted models of commercial off the shelf products and include known and unknown biases; the beneficence of using energy-intensive AI algorithms for 'entertainment' and speculative research; fears about replacement of humans or dumbing down comedy; and uncertainty about how the results of studies into machine comedy will be used (see Lucie 2021, 319, Mirowski et al. 2024, Mirowski et al. 2025). However, as mentioned, AI is already being used by comedians, and this usage is not likely to decrease. Understanding its impact and limits and speculating on potential future affordances that are specific to the field is timely and necessary. Considering the ways in which AI might engage with and learn from comedy performance offers the distinct benefits of picking up on aspects of audience experience that humans cannot, such as perceiving patterns in delivery that are imperceptible to the human sensorium. Quantitative measures of audience experience such as self-reporting (Tung Au et al. in Reason et al. 2022, 326) have already introduced data into the field of audience studies that present alternatives to the individual, expert critic with their particular taste. An AI trained to read affective responses could provide a different approach to existing quantitative methods for gauging real-time audience engagement, such as cognitive neuroscience, which can be used to 'measure what spectators do and feel' through body movements, behaviors, and physiology (Millman et al. in Reason et al. 2022, 293). By analysing this and other data such as human speech (Tang et al. 2025), the use of AI tools offers to expand the analytical tools available to the field of audience studies, potentially expanding what the field understands as its source data in the process. We agree with L.S. Merritt Millman, Daniel C. Richardson, and Guido Orgs (Reason et al. 2022, 293) that quantitative methodologies complement established qualitative methodologies rather than replacing them. Our purpose here is collaboration and augmentation, and to consider what could be done with new kinds of performance data and analysis. To this end, the paper begins by introducing the more familiar question of AI's strengths and weaknesses for writing comedy, which several comedians and scholars are already exploring. The rest of the paper is structured to speculatively address the following questions: What are AI's limitations and affordances for qualitative analysis of comedy audiences? What would it take for an AI to read the room live? And, if AI could read the room, could it then learn to perform believable comedy?

Writing comedy with AI: limitations and affordances

Some stand-up comedians have been quick to embrace the recent wave of generative AI, with some even incorporating it into their acts, thereby provoking questions about what it takes and means to be funny in a live performance. Often it is AI's failures or inadequacies that become the focus of these acts, raising further questions about

how to disentangle laughing at AI from laughing at its material. For example, in 2024 Canadian comedian Anesti Danelis collaborated with ChatGPT to write his musical comedy show, *Artificially Intelligent*, which he performed at the Edinburgh Festival Fringe and the Toronto Fringe Theatre Festival (Lawton 2024). In interviews Danelis distinguished between what ChatGPT was good at—namely, devising the structure of a show—and its 'hilariously terrible ideas', which became the show's content (Lawton 2024, CTV Your Morning 2024). The hilariously terrible ideas involved reducing the data Danelis entered about himself to stereotypes about his Greek immigrant family background or his millennial generational identity (BBC Scotland 2024). In clips from the show uploaded to TikTok, Danelis speaks to a glowing blue orb that rests on a chair. The onstage orb, which speaks with a male-sounding voice, is intended to represent ChatGPT, which displays an onscreen blue orb when users interact with it in Advanced Voice Mode (Fitzpatrick 2024). When Danelis objects to the platform's stereotyping, the orb replies, 'From what I gathered on the internet this is who you are' (BBC Scotland 2024). This points to the limitations of what and how ChatGPT 'reads' in its construction of Danelis' comic persona. Danelis noted that to help the platform execute his instruction to write him a song in his own voice, he 'kept prompting it with more information on my life' (CTV Your Morning 2024). 'I would ask it to write jokes in the style of me and it didn't understand me', Danelis added in another interview (BBC Scotland 2024). Despite feeding in more data about himself, Danelis could not get the platform to produce jokes in his voice. It is likely that in the show, Danelis has not transcribed ChatGPT's responses to his prompts verbatim but instead has drawn inspiration from its reductive construction of persona in order to make a point about its limitations as a reader. This also points to how the platform's capabilities as a reader are concomitant with its limitations and affordances as a writer.

AI potentially disrupts Bergson's categories, not only because definitions of the human have expanded, but also because AI brings adaptability and learning into the technological. But at present AI's adaptability and learning are limited by medium, generalist bias, and blindness to the context of delivery. Danelis' experience of collaborating with ChatGPT speaks to many of the concerns articulated by other comedians in a 2024 study by Piotr Wojciech Mirowski et al. Noting that the complexity and human skills involved in comedy expose some of the limitations of large language models (LLMs), the study's authors interviewed twenty comedians about their experience using AI for comedy writing. The study explored several hypothesised limitations of LLMs for writing comedy: 'stereotypes, inability to distinguish comedic offensiveness from harmful speech, cultural erasure and homogeneisation [*sic*] of content' (Mirowski et al. 2024, 3). Participants largely confirmed these limitations, though they also, like Danelis, affirmed AI's usefulness for getting a structure or an improvable first draft in place (Mirowski et al. 2024, 5). Participants further noted that AI has some fundamental limitations in relation to comedy's conventional values and mode of delivery: these limitations include the 'inability to pull from personal experience, lack of perspective, and lack of context and situational awareness' (Mirowski et al. 2024, 6). While the comedians in the study were rightly circumspect about AI's usefulness to their craft, it is precisely the fact that comedy encompasses a lot of things LLMs cannot currently do or understand that places this genre at the frontier of thinking about them.

Reading the digital room: AI and qualitative audience analysis

The lack of context and situational awareness especially means that AI cannot read the live room in the traditional sense. But it can already read some limited versions of the expanded, mediatised room. For example, although the free version of ChatGPT cannot directly process video files, it can describe and analyse the audience of comedy performances about which it has sufficient data. We asked the platform to tell us about the audience of Hannah Gadsby's *Nanette* (2017, 2018 on Netflix), a show that played a prominent role in international debates about comedy, misogyny, and homophobia during the #MeToo movement against sexual violence. We selected this show for several reasons: (1) its prominence in these debates means there is a lot of public discourse about it, which gives ChatGPT ample data (of a particular kind) to work with; (2) its status as an influential example of what some people call 'anti-comedy' or 'post-comedy' (Giuffre 2021, 29) brings into relief the sorts of conventions and departures from convention the AI can pick up on; and (3) one of this article's co-authors has published considerably on *Nanette* (Balkin 2020; Balkin 2023; Balkin and Mierowsky 2024), and her familiarity with the show and its manipulation of audience response positions her well to judge the strengths and weaknesses of an AI's audience analysis. This example is therefore chosen for its exceptionality as much as its representativeness; as the following discussion shows, both of these aspects of it (the ways it is and is not conventional comedy) come through in the AI analysis. This is relevant for thinking through the limitations and affordances of a platform whose perceived and actual weaknesses include generalised bias. Thus, the analysis in this section is necessarily one key example (rather than a comparative survey) of AI's ability to read the expanded room.

ChatGPT gave an efficient breakdown of *Nanette*'s key audience groups: 'LGBTQ+ Community', 'Feminists & Social Justice Advocates', 'Comedy Fans (Especially Those Who Enjoy Thought-Provoking Stand-Up)', 'Critics & Intellectuals', 'Survivors of Trauma', and 'People Curious About New Forms of Storytelling'. It ended by asking whether we would 'like insight into how different demographics reacted to *Nanette*?' ('What Can You Tell Me About Nanette's Audience?' 2025). The platform recognised reception of Gadsby's comedy special as 'divisive' and could break down the responses accurately by demographic group. It could, in other words, read the expanded digital 'room' insofar as that room is the internet's extensive body of textual commentary on *Nanette*, which was an international phenomenon that inspired a great deal of controversy and commentary (see Balkin 2020).

What ChatGPT could not do was analyse audience responses such as laughter, applause, and tense silence in the recorded Netflix version of *Nanette*, which notably eschewed the conventional practice of cutaway shots to keep the focus on Gadsby and maintain tension in the filmed performance (Giuffre 2021, 36). It was upfront about the kinds of data it could read: when we asked, 'Can you analyse the audience responses in the Netflix version of Nanette?' it replied, 'I can't analyse audience responses directly from the Netflix version of *Nanette*, but I can discuss how audiences reacted based on reviews, social media discussions, and critical analysis'. It then gave several excellent suggestions for how to achieve 'a more data-driven breakdown (e.g. audience laughter, applause, silence, or specific moments of engagement)', including

observing crowd reactions during key moments, comparing with more traditional stand-up specials to see the differences, and looking at online discussions and reviews for evidence ('Can you analyse the audience responses in the Netflix version of Nanette?' 2025). This is the same kind of advice researchers in theatre and performance studies would give to students looking for evidence of audience response to a performance they had not seen live. Applause, laughter, silence, and other reactions constitute the audience's own performance, which co-creates the event (Heim 2). This is especially true in stand-up, which emerges from the vaudeville tradition of an outwardly responsive audience, and where those responses are openly acknowledged by the comedian as part of the performance situation (Balkin and Mierowsky 2024, 40; Double 2014, 20; Jenkins 1992, 31-32). Thus, although the free version of ChatGPT cannot currently read the recorded 'room' of a comedy performance in the sense of analysing the audience's gestures, vocalisations, and silences, it recognises the nature of the data it cannot currently access and what should theoretically be done with it by a student or scholar of theatre and performance.

While Mirowski et al. (2024, 7) rightly note that the text-only medium of LLMs limits their usefulness in a genre that relies on delivery and other aspects of performance, this is changing. At the time of writing ChatGPT has introduced an 'Advanced Data Analysis' function for paying Premium members, which supports audio, video, and image files.[4] Media reporting on this new ability to process video inputs notes that AI can now analyse 'real-time visual insights' and 'give context-rich answers' (Kahekashan 2024). These functions may support different modes of audience research in future. Gemini AI, a free chatbot developed by Google, which also owns YouTube, can already analyse video when users paste in a YouTube video URL.[5] Gemini is one of several multimodal large language models (MLLMs), which can operate across modalities including language, visual, auditory, and sensory data (Fei et al. 2024). We asked Gemini, 'Can you watch this video and analyse the audience reactions?' and provided the URL of a publicly available nine-minute clip from *Nanette* (Netflix Is a Joke 2021).[6] The video is from a part of the show in which Gadsby speaks angrily and details a traumatic experience of sexual assault; it does not contain many jokes. As noted earlier, *Nanette* also eschews the conventional practice of cutaway shots that show audience members' facial expressions, limiting Gemini's data sources in the clip. Gemini responded, 'The video features Hannah Gadsby sharing a powerful message about the experiences of women and the need for men to acknowledge their role in creating a safer world. Here's an analysis of the audience reactions, based on the content of the video'. By the content of the video Gemini appeared to mean Gadsby's spoken language; its analysis of audience responses did not refer to data about the specific Netflix audience whose reactions are audible (but not visible) in the clip. Moreover, Gemini framed its analyses of audience responses as speculative; for example: '**Admiration and Inspiration:** The emphasis on resilience and strength, with Gadsby refusing to be defined as a victim and highlighting the strength of a 'broken woman who has rebuilt herself', likely inspired and moved the audience [06:25], [08:33]' ('Can you watch this video and analyse the audience reactions?' 2025, bolding and timestamps in original). Each of Gemini's listed audience reactions, which also included 'Anger and Frustration', 'Empathy and Connection', 'Agreement and Empowerment', and 'Understanding and Validation', included the word 'likely' in this

way. The timestamps refer to relevant moments in Gadsby's speech, but not to particular moments in the audience's response. The audience Gemini described was therefore a general rather than specific one.

However, when we followed up to ask, 'Can you analyse specific audience reactions in the video such as laughter and applause?', Gemini provided four different timestamps for audience reactions in the video, but did not comment on them. The timestamps seemed not to be entirely accurate: although Gemini noted an audience reaction at 3:34, there is no laughter or applause then; this is a moment in the show when Gadsby speaks angrily. Revealingly, when we followed up again to ask, 'Can you comment on the meaning of these moments of laughter and applause?' (2025), Gemini replied:

> **[00:03:34]:** At this timestamp, Hannah Gadsby reflects on the painful experiences she endured, stating that they stemmed from shame and hatred directed towards her identity. While the tool output doesn't explicitly mention laughter or applause here, the gravity of the statement suggests that the audience reaction might have been one of silence or perhaps a more subdued acknowledgment of the seriousness of her experience.

This made us reconsider our prior question, which asked Gemini to consider specific audience reactions *such as* laughter and applause; that is, we had not limited the question to those responses. Gemini, then, had extrapolated what might be described as a passive audience response—tense silence—even though it was not specifically prompted to note that reaction. Gemini was also able to identify emotional shifts across the arc of the clip; it explained another timestamped moment, 00:04:27, as follows:, as follows:

> Immediately after saying, 'That is your last joke', there is likely laughter from the audience. This moment signifies a shift in the tone of the performance, where Gadsby deliberately moves away from traditional comedy. The laughter could be a reaction to the unexpectedness of this statement in a comedy setting, or perhaps a slightly nervous acknowledgment of the change.

The word 'likely' suggests Gemini was unable to say definitively that what it heard is laughter; in this instance it was correct. Gemini was also able to identify the affective conventions of a comedy setting and some of the ways *Nanette* works against them when it refuses to release the audience's tension through laughter (Balkin 2020, 79). Thus, the AI can interpret affective cues such as silence, laughter, and applause, and can make claims about their meaning with reference to the content of the specific recorded performance and more general comic conventions.[7]

This section models how AI can read the expanded digital room in some of the ways that theatre and performance scholars analyse audiences. As noted, Gadsby's *Nanette* is only one example and is in some respects exceptional rather than representative; a comparative analysis would provide more generalisable claims. At the same time, several of the ways in which Gadsby's show is unconventional—its inclusion of 'serious' or tense sections and its questioning of comedy's purpose, for example—have become comic conventions in their own right, suggesting the broader applicability of this analysis.[8] With the right prompting and access to relevant datasets, AI is on the verge of detecting (and in some instances can already identify and convincingly explain) seemingly ephemeral and affective aspects of performance, such

as tense silence. The next section asks how an AI comedian might analyse audience response and affect in live stand-up, where reading the room is a precursor to responding to the audience and maintaining the performer's license to speak.

From text to performance: AI learning from affective atmospheres in performance

AI and robotics comedy research has been primarily focused on the production of humour in language and is increasingly branching into live performance as the 'natural' site for its evaluation, 'in front of audiences sharing either physical or online spaces, and under real-time constraints' (Mirowski et al. 2025, 1). Translating AI comedy texts produced by LLMs into embodied performances is a key site of creativity, with experimentation ranging from desktop computers visible on stage, to humanoid robots, to headphone-wearing human avatars, to the blue orb mentioned above. A significant majority of the field is preoccupied with stand-up comedy, with a trope of anthropomorphic robots performing LLM-generated stand-up comedy. Roboticists' interest in stand-up may stem from the form's reachability combined with the perceived benefits of a 'repeatable [theatre] setting in which research variables can be tested, the ability to instrument [*sic*] and control the environment, and the many participants in the audience from whom one can collect data from simultaneously' (Swaminathan et al. 2021, 1262). While this vein of work has advanced understanding of the production of comic language, with some examination of timing and body language, significant foundational elements of comedy have been overlooked, including: performance, presence, and the ability to 'read the room'; staging, scenography and environmental affordances; combinations of human and machine performers; and, non-verbal, visual, and physical comedy. Put plainly, AI comedy has been trained on a corpus of literature and jokes, but not of performances, how to tell jokes, or how to be funny without words.

The translation of scripted jokes into live comedy amounts to a shift from a written language ontology to a performance ontology, understood in theatre practice as the page-to-stage process undertaken in rehearsal: translating written text into a performance text, through the embodied presence of the performer. Central to the ontology of live comedy is that each performance is enacted anew in a different way, responsive to the specific context of place, audience, and moment. Hence 'reading the room' is a term for the performer's ability to create comedy by adapting rehearsed material live within the affordances they perceive during each moment of performance by inflecting their timing, intonation, gesture, facial expression, emphasis, and even choice of material. Reading the room is one of the skills Beck Krefting identifies as necessary to 'manage audience responses, take an ongoing emotional inventory of audience, and develop content that will effectively elicit laughter across communities adhering to affective conventions of the art form' (2024, 4). Such presence within the live unfolding production of a comedy performance is a skill acquired through practise; it is learned by making and participating in comedy and is irreducible to scripted jokes alone.

While some work has commenced on a robot or AI system's 'ability to read the room and work the crowd' (Gray, Myers & Fitter 2020, 12165), high technical

complexity and reliance on multiple modes of sensing have limited progress. However, the problem itself has not been adequately defined. How do human comedians read a room? Is it possible to emulate this sensory-cognitive ability in machines, and if so, how? Might similar human-level 'readings' be achieved in other ways by providing machines with combinations of sensors and training different from human sensing-intuiting? Is the intention to make AI and robots funny like humans, or for them to produce new hybrid and more-than-human forms of comedy? A common aim, repeated in robot and AI comedy studies, is that making machines emulate human comedy and humour will increase their acceptance and uptake, but is that the best approach for stand-up, which is traditionally bound in the idiosyncrasies of the individual performer? Might a pluralised, non-normative approach to discovering 'new species' (Ullrich et al. 2020, 1) or genres of comedy be more appropriate?

A human comedian's ability to read the room might first be based on their ability to attune to their environment through a variety of sensory means, including sight, vibration, proprioception, and especially their ability to sense the audience through sound. The lighting of many stages makes seeing the whole auditorium, even in small venues, difficult. Movement in the audience is often overlooked by necessity and the physical distance of some audience members means that any meaningful 'reading' of their facial expression is impossible. The asymmetrical ocular relationship between stage and auditorium is such that while everything the performer does physically is hyper-visible and lit by powerful theatre lights, the audience is less visible: in larger theatre venues most audience members are anonymous in a darkened auditorium, while in more intimate venues and club settings the comedian may be able to see most or all audience members, though the performer is still the ocular focus. The amplified nature of standup's iconic microphone and the loudness of the audience means that the work is often experienced as vibrational noise in the body in addition to the recognisable, enculturated sounds of the ear. The interplay for dominance of the acoustic space unfolds through the rhythm of the comedian's amplified voice and the audience's individual and collective capacity to create and withhold sound. Hence the reading of 'reading the room' is less about understanding the audience as a visual text, and more about an audial appreciation of qualities of a sound texture. It is more like taking an atmospheric reading of a room, like a weather balloon, than reading it like a book.

Early definitions of the term 'audience' denoted hearing (including attending a court hearing and having the opportunity to be heard) rather than seeing (*Oxford English Dictionary*, 'audience (*n.*)', March 2025). Shakespeare's audiences 'talked about going to *hear* a play and going to *see* one in equal measure' (Folger Shakespeare Library 2021). The *Oxford English Dictionary*'s earliest (1728) definition of 'auditorium', 'that Part of the Church where the *Auditores* stood to hear, and be instructed' ('auditorium (*n.*)', March 2025), likewise denotes a space for listening or auditing. The invention of gaslight, increasing pictorialism, and emphasis on spectacle in nineteenth-century theatre all contributed to the dominance of visuality in modern theatre.[9] Stand-up comedians, whose work is often visually simple, in some respects still occupy the intense sound world of auditoria, where their perception of the flow, timing, success, and ultimately the possibility for them to work the crowd is afforded in large part by their reading, their intense listening-intuiting-presence, of the room.

The range of possible readings is as diverse as reactions to their material, which vary from one performance to the next even as they tend to fall into patterns which skilled comedians can predict and for which they can prepare seemingly off-the-cuff responses. Considering mimicking this skill through machine listening reveals some of its complexities and shortcomings. To begin scoping the challenge, consider the qualities and range of laughter, its distribution over the audience, its abruptness, endings, volume, beginnings, and duration. Contrast this with the same qualities for silence, or the withheld noises of audiences. And then the range of other non-laughter responses, including applause, cheers, groans, gasps, and sighs. Then, the three-dimensionality of mapping sections of the audience to know who is responding to what and where. Hecklers, disturbances, and solo ruptures stepping out of the homogenous chorus of a mass audience also draw the attention of seasoned comedians as they offer new affordances within the present unfolding of a set.

To read and work a room, a machine stand-up comedian would require the sensory apparatus to perceive its affective atmosphere, the sonic topography of an audience, and the intuitive and cognitive capacity to map its trajectory over the course of a set, measure responses to jokes, establish a rhythmic rapport with an audience, adapt both the content and performance of jokes based on their readings, and develop a capacity to perceive creative affordances in disruptions and heckling. This seems a tall order for LLM artificial intelligence trained primarily on corpora of written material and not live performance. To learn these skills as humans do, machine comedians need to train on stage with live audiences within the affective atmosphere of performance.

However, many machine learning techniques observe and gain insights into human behaviour using sensors and combinations of data in modes beyond the human body's ability to sense or compute. For example, there has been significant research on detecting mental health by combining a range of observations from smartphone use without listening to conversations. Research that connects the timing of frequency of phone activations (such as during the night), distances travelled each day, the number and frequency of conversations occurring around the smartphone (not the contents), and time spent in certain apps, may be applicable to comedy (see Khoo et al. 2024 for an overview). Such multi-sensor, passive detection might underpin machine attempts to perceive affective atmospheres.

While a human performer is limited to seeing and hearing based on the location of their head, a machine comedian's sensory system might be distributed around an auditorium to enable it to read the room in unprecedented ways. While a human performer may use sound to intuit the affordance of the affective atmosphere of a given moment during a set, cameras could track the facial responses of every single audience member. Similarly, distributed microphones may reveal more granular information about specific regions of an audience's affective topography. Most audience members would likely carry a smartphone, and while many human comedians confiscate smartphones as a condition of entry to a show, a machine comedian might instead recruit the smartphone's sensors to its distributed room reading endeavour. Smartphones and smartwatches might reveal basic information about their owner, from when or where they check the device during the performance to far more personal bio-data, including their heart rate, movement, and the volume of their personal responses. Such multiple modes of data on each individual audience member, their

immediate area, and the audience as a whole would provide a wealth of data to machine comedians who could compare the loudness of a general audial response with fine-tuned changes in somatic state to truly machine their working of a room. A further fine reading of the room could come from incorporating social media profile data into sets, including, for example, shows watched on Netflix and YouTube of specific comedians and their styles. Such data could be extrapolated from ticket purchasing information cross-referenced with a seat number and could be further specified if personal smart devices were accessible to the machine comedian. How, for example, has my audience responded to my material in the past? What are they interested in? What are their jobs, political beliefs, and relationship status? These are not new questions for performers—indeed, 'What do you do for a living?' is one of the great cliches of crowd work—but combining in-the-moment sensory data with these other kinds of datasets suggests the expanded kinds of reading an AI comedian could undertake.[10]

What is the purpose of studying AI in comedy beyond the instrumental desire of computer scientists, roboticists, and programmers to incorporate it into more effective and appealing (read tolerable) interfaces? The computer linguistics community have studied humour for decades and consider it to be an 'AI-complete problem' because of its reliance on human cognitive abilities such as memory recall, semantic integration, world knowledge, and cultural context, making it as challenging as achieving artificial general intelligence (Winters 2021, 3). Despite significant progress in machine humour detection and generation, 'achieving a true and full computational sense of humor is still far off' (Winters 2021, 8). The potential of machines, therefore, to create 'comedy as an intentionally structured cultural product' (Meany 2016, 169) appears remote. Recent surveys of the 'comedy community, from actors and audiences to reviewers and journalists writing about comedy' suggest that they share this 'skeptical view about AI's comedic potential' (Mirowski et al. 2025, 1). Yet, despite the overwhelming challenge, interest in machines' ability to create and perform comedy is widespread in academic research, industry, arts practice, and popular culture. This is because comedy is understood to be a distinctly human trait intertwined with social identity, cultural context, place, and personhood. However, even if artificial general intelligence and life-like humanoid robots were achieved, a significant 'anchoring effect' would curtail acceptance of their potential and prevent people from believing machines can transcend machine-ness (Ullrich et al. 2020, 2). A study of hundreds of participants suggested that human prejudice against anything named a robot 'can hardly be overcome' (*ibid.*), presenting a paradox: people desire machine humour but seem to have unconsciously foreclosed on ever accepting it.

Machine performances of comedy often fail because we don't believe their authenticity, while unintended humour emerges effortlessly as a by-product of their all too believable failure. It is not, therefore, merely a question of intention, of wanting robots to be funny, but of transcending a belief that because a performer is a robot it cannot be deliberately, authentically, and legitimately funny, the way a human can. By inventing machines in our own image and setting them up to fail at reproducing human humour, we miss that it may be 'more promising to understand robots as their own 'species'' (Ullrich et al. 2020, 1) and, therefore, to seek out a new genre of comedy made by and with machines. Recent developments in post- and more-than-human design practices (Bleeker and Rozendaal 2021; Gemeinboeck 2021; Gemeinboeck and

Saunders 2023; Giaccardi and Redström 2020; Lee et al. 2025) suggest that researching 'the fundamental differences between humans and robots as a starting point for the creative development' (Rozendaal et al. 2025, 1) of new kinds of comedy performance might be more productive than training machines to perform human comedy. The question is: are there new genres of comedy that are 'more-than-human', created by and with machines, and are they recognisably funny to people? Further, any sufficiently alien or new form of comedy will beg the question: can the room read AI?

Social and cultural theorists have developed several (sometimes overlapping) paradigms that theatre and performance scholars have used to understand human and nonhuman collaboration in performance. Prominent examples include Bill Brown's 'thing theory' (2001), which emphasises a blurred boundary between subjects and objects; Bruno Latour's actor-network theory (2007), which posits relational networks of human and nonhuman 'actants'; new materialism more broadly, which emphasises the agency of nonhuman matter; and posthumanism, which challenges anthropocentric modes of thinking and being. However, with the focus of comedy on producing all-too-human laughter in human audiences, it is probably fairer to sidestep post- and more-than-human stand-up, to instead favour a term that emphasises machine enhancement while acknowledging the human-centred aim of comedy. While the definition of the human has expanded, and what humans do in collaboration with AI is becoming part of comedy for that expanded human audience, comedy practice is still focused on making people laugh. If the room cannot read the AI, its comedy cannot succeed.

Toward augmentation and collaboration

Our article shows that AI is, in some ways, on the cusp of recognising the kinds of affective atmospheres and cues that comprise audience response to stand-up. It can, within the limits of the datasets available to it, analyse comedy audiences effectively. This suggests MLLM AI might afford audience studies new modes and methods of analysis, albeit with ethically and methodologically problematic opacities around its datasets. What is less clear is whether a time will come when AI can be deployed to create comedy that humans find funny by adapting prepared material to meet the immediate context of the 'room' in which its performance occurs. This is what is conventionally at stake in reading the room: the comedian's skilled analysis of the audience's reactions licenses the performer to continue speaking. This comic license is the agreement or contract between performer and audience. As Sarah Balkin and Marc Mierowsky (2024) discuss, this contract is conventionally predicated on the comedian's persona, their performed self. This performed self provides the (real or fictionalised) personal experience and perspective that the comedians in Mirowski et al.'s study (2024, 6) see as missing from comedy written by LLMs. So, having the wherewithal to read the room alone would not grant the machine performer comic licence: they will also be required to pass as a believable persona. But need this persona be recognisably human?

At present, stand-up remains a largely human art, though comedians have a long tradition of collaborating with nonhuman objects and technologies (microphones, props, sound and lighting effects, PowerPoint slides, digital distribution platforms). We suggest that currently, AI's contributions to stand-up should be understood as part of this tradition of collaboration between human performers and nonhuman aspects of production,

rather than as a new species or genre of stand-up. As comedy practice and innovations in machine performance are both currently purposed toward entertaining people, a radically new genre is unlikely to emerge. Machine comedy for machines seems a tall order, but is not unthinkable. Imagine, for example, LLMs taking opposing positions to rehearse comedy into being. Such adversarial models are common in AI and machine learning algorithms. However, in the short term, machine intelligence's entrance into comedy is likely to remain in service of mimicking what humans find funny.[11]

Hence, if digital cohabitation describes the world humans and LLMs increasingly occupy together, digital collaboration might best describe the current relationship between (M)LLMs and the comedians who deploy them in stand-up. In examples like Danelis's show the AI becomes a kind of unintentional straight man who cannot adequately read Danelis or the live room; Danelis deploys these failures in a performance that is funny to humans because the human performer repeatedly points to the ways in which the machine performer is acting like a machine. Seen in this way, we could include AI's tendency toward generalist bias and stereotyping among the Bergsonian rigidities to which comedy offers a social corrective. For Bergson, the question of why we laugh is tied to this social purpose of facilitating human adaptivity. But, although it is clear from our AI-assisted analysis of Gadsby's *Nanette* that MLLMs can detect some of comedy's affective conventions and social meanings, modes of interpretation predicated on producing social meaning may ask the wrong sorts of questions for training MLLMs in performance.

The habit that produces laughter in the body from a sensual attunement to an affective atmosphere, often occurs prior to consciously knowing *why* something is funny. The ambiguity of not knowing why we laugh but find ourselves doing so anyway might be ascribed to an affective contagion.[12] Habitual response to affective atmospheres proceeds through physical, though often impalpable, cues to the sensorium stimulating and updating an interior mirror of the immediate environment and the perceived situation unfolding in a performance. Considered in this way, habits help us consider the parallels between human and AI learning (Chun 2017, 396; Walton 2021, 290). This raises the question: does the machine need to know *why* we laugh, or *how* to make us laugh? The evolution of machine learning would suggest the latter, that there is no need for AI to understand why something is funny as, 'when the same algorithms are audience to billions of performances within the same given context, *why* is not required: statistics reveal the habits of large human populations and make behavior predictable' (Walton 2021 294). Similarly, we posit that comedians and audiences do not always know why we laugh, but have a tacit sense of humour that provides a how. The desire to understand why is reminiscent of a foundational question of whether computational processes can originate anything new and therefore surprise us, rehearsed by Alan Turing as 'Lady Lovelace's Objection' (1950, 433). An MLLM, 'as an analytical engine has, to paraphrase Lovelace, no pretensions to originate anything, it merely must perform its computations to predict what is likely to happen next, before what happens next happens' (Walton 2021, 294). Similarly, an experienced stand-up navigates the affective atmosphere of a live performance to produce comedy that sometimes exceeds their pre-rehearsed material and conscious understanding to meet the moment. This is why MLLMs must be trained in performance and not merely by reading scripts, accounts, and watching

videos, as they must attend not only to the non-verbal aspects of comedy underlying stand-up, but also to the ways in which comedians refine their acts in response to audiences across a run. Attunement to affective atmospheres—reading the room—exceeds language. But what would such training look like?

For this we could turn to Philippe Gaulier's 'pedagogy of spectatorship' to consider appropriate training for AI stand-ups in performance to ask: what would happen if the only feedback we trained AI on was laughter and other non-verbal responses? In Gaulier's regime, trainees 'develop skills in spectating the audience' (an alternative 'reading the room') and learn to 'claim agency over their own ridiculousness' by responding to what makes the audience laugh in real time (Amsden 2016, 15). Enabling AI to seek out and claim agency over its own ridiculousness, that is, to become aware of how people perceive its machine-like performances and failures to perform human persona, and to potentially repurpose these discoveries to deliberate comic effect, offers a potential technique for discovering comedy through an immanent *via* negativa training paradigm. Such a paradigm would be especially appropriate to stand-up, a genre that 'maximizes the experience of failure' by requiring its practitioners to rehearse before a live audience (Lintott 2020, 214). This all-too-human training paradigm, in turn, raises questions about the degree to which pattern recognition is separable from meaning and social purpose.

While this paper explored the emerging capabilities and limitations of AI in stand-up comedy, particularly focusing on the complex skill of 'reading the room', there is still much to do. AI has made strides in analysing comedy audiences and generating humorous text, yet we argued that translating this to live performance requires a more nuanced understanding of timing, delivery, and audience interaction, and we speculated on as-yet distant techniques for advanced machine 'readings' of rooms. We proposed that rather than attempting to replicate human comedy, research could explore new 'more-than-human' forms of machine-generated humour, while remaining cognisant that the target of comedy remains a human audience. Currently, AI's role in stand-up is best understood as augmenting human performers rather than replacing them. We suggested that training AI in actual performance settings, not just text generation or video analysis, is crucial for developing skills like real-time audience analysis and adaptation. Drawing parallels between human and machine learning in comedy, we suggested that both rely on pattern recognition and habit formation to arrive at something funny within the affective moment of performance. We speculated on a future training paradigm for AI in comedy that teaches machines to recognise and leverage their own ridiculousness. Ultimately, AI stand-ups require both an ability to perform comedy to meet the affective atmosphere of live performance and the ability to perform a believable persona to acquire and maintain comic license, which permits them to speak. It remains to be seen if either of these two fundamental requirements will ever be achieved by machines.

Notes

1. We thank Wally Smith for his insights as we developed this article.

2. AI also has many applications for humour more broadly. For a succinct overview of these applications (as of 2023), see Nickl, Benjamin, and Christopher John Muller. 2023. "The Joke's on Us – How AI Is Replicating Our Laughter Online." *The Conversation*. September

21, 2023. https://theconversation.com/the-jokes-on-us-how-big-tech-is-replicating-our-laughter-online-206191.

3. Rejecting the active/passive binary is now a commonplace in scholarly accounts of theatre audiences; for some key readings on this topic see Rancière, Jacques. 2011. *The Emancipated Spectator*. Translated by Gregory Elliott. London: Verso; Freshwater, Helen. 2009. *Theatre & Audience*. Theatre&. Houndmills, Basingstoke, Hampshire, UK; New York: Palgrave Macmillan; and Heim, Caroline. 2016. *Audience as Performer: The Changing Role of Theatre Audiences in the Twenty-First Century*. London and New York: Routledge. Often accounts of an audience's conventional behavior are grounded in historical and embodied conditions of spectatorship, such as the modernist darkening of the auditorium. See for example Kershaw, Baz. 2001. "Oh for Unruly Audiences! Or, Patterns of Participation in Twentieth-Century Theatre." *Modern Drama*, no. 2, 133–54; or Julia Walker's (2022, 229) account of how the introduction of air conditioning reshaped the experience of theatre and film spectatorship.
4. This function was formerly called "Code Interpreter," which is worth mentioning in light of the acknowledged performativity of code (Metzger 2021, xiii).
5. We speculate that Gemini is able to offer video analysis as part of the free AI service because Google owns YouTube; these videos therefore may comprise part of Gemini's readable data repository. Gemini's training data comprises "publicly available sources" ("What is Gemini" n.d.), but the company is not transparent about what this includes.
6. The availability of clips from *Nanette* on YouTube is another reason for our selection of this example.
7. Gemini marked another seemingly inaccurate instance of audience response at 9:11, one second before the end of the YouTube video, when it shows an ad for Netflix subscription that is not part of Gadsby's show. Since the diegetic audience is not present for this final part of the clip, Gemini's reference to "applause and cheering at the end of the video" when we asked it to comment on the meaning of this moment is a summation of the audience's final response rather than a reflection on a response happening at the time indicated.
8. See Balkin 2023 and Balkin and Mierowsky 2024 on parallels among Gadsby, Dave Chappelle, and Hasan Minhaj's Netflix specials and they ways they use seriousness, sincerity, and humorlessness.
9. See Crary 2005. On the effects of historical advances in stage lighting see for example Penzel, Frederick. 1978. *Theatre Lighting before Electricity*. 1st ed. Middletown, Conn: Wesleyan University Press; and Baugh, Christopher. 2013. *Theatre, Performance and Technology: The Development and Transformation of Scenography*. Theatre and Performance Practices. Palgrave Macmillan.
10. This dystopian data-surveillance vision of the potential to tailor a performance to specific audience tastes and preferences is more likely to evolve outside of live auditoria, where complex groups of audiences gather. Instead, channels accumulating and making decisions about viewing data, like Netflix and YouTube, might well offer fully bespoke comedy material tailored to individual viewers' desires and humour. The more you watch, the funnier it becomes. Or would it?.
11. In this sense, like other contemporary AI pursuits, the current state of the art is what Chun describes as homophilic, "love as love of the same", of "discriminating patterns" in human preferences: at once working into them while simultaneously reinscribing them into human preferences (Chun 2018, 62; Chun 2021, C2).
12. On emotional contagion in theatre practice and discourse see Fischer-Lichte, Erika. "How to Conceptualize Emotional Contagion in Performance." *Habitus in Habitat I: Emotion and Motion* 1 (2010): 25-40.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Notes on contributors

Dr Sarah Balkin is a Senior Lecturer in English and Theatre Studies at the University of Melbourne. Her current research focuses on the historical emergence of deadpan performance styles (1830-1930) and changing relationships between performers and audiences in contemporary comedy. She is the author of *Spectral Characters: Genre and Materiality on the Modern Stage* (Michigan 2019) and the co-author (with Marc Mierowsky) of *Comedy and Controversy: Scripting Public Speech* (Cambridge UP 2024). She has published widely on comedy, audiences, genre, theatre, and literature from the nineteenth century to the present. She is the guest editor of this special issue on The Comic Audience.

Dr Robert Walton is an artist and director recognised with multiple awards for his work in theatre, screen, installation, writing, interactive art, and research. He is the Dean's Research Fellow at The University of Melbourne's Faculty of Fine Arts and Music. In this role, he leads the development of performances and artworks that explore the creative potential of ancient and emerging technologies, including artificial intelligence/automated intuition, theatre, virtual holograms, swarm robotics, standing stones, engineered bacterial bioluminescence, algorithmic ritual, MR/XR, storytelling, building information model data, ambient computing, expanded photography, virtual production, and fire.

ORCID

Sarah Balkin http://orcid.org/0000-0002-1430-9063
Robert Ellis Walton http://orcid.org/0000-0003-2857-3372

References

Amsden, Lucy. 2016. "When They Laugh Your Clown is Coming': Learning to be Ridiculous in Philippe Gaulier's Pedagogy of Spectatorship." *Theatre, Dance and Performance Training* 7 (1): 4–16. https://doi.org/10.1080/19443927.2015.1133445.

Balkin, Sarah. 2020. "The Killjoy Comedian: Hannah Gadsby's *Nanette*." *Theatre Research International* 45 (1): 72–85. https://doi.org/10.1017/S0307883319000592.

Balkin, Sarah. Mar 2023. "On Quitting: Dave Chappelle's *The Closer* and Hannah Gadsby's *Nanette*." *TDR: The Drama Review* 67 (1): 149–166. https://doi.org/10.1017/S1054204322000958.

Balkin, Sarah, and Marc Mierowsky. 2024. *Comedy and Controversy: Scripting Public Speech*. Cambridge: Cambridge University Press. https://doi.org/10.1017/9781009396820.

BBC Scotland. 2024. "Can Artificial Intelligence Write Decent Jokes?." TikTok. August 17, 2024https://www.tiktok.com/@bbc_scotland/video/7403753975431941408.

Bergson, Henri. 2003. "Laughter: An Essay on the Meaning of the Comic." Translated by Cloudesley Brereton and Fred Rothwell. Project Gutenberg http://intersci.ss.uci.edu/wiki/eBooks/BOOKS/Bergson/Laughter%20Bergson.pdf.

Bleeker, Maaike, and Marco Rozendaal. 2021. "Dramaturgy of Devices: Theatre as Perspective on the Design of Smart Objects." In *Designing Smart Objects in Everyday Life: Intelligencies, Agencies, Ecologies*, edited by Marco Rozendaal, Betti Marenko and William Odom, 43–56. London: Bloomsbury Visual Arts.;

"Can You Analyse the Audience Responses in the Netflix Version of Nanette?" 2025. *ChatGPT*. 14 February. https://chatgpt.com/g/g-TDjzPM7vk-video-summary-and-analysis/c/67aec5e1-ee1c-8012-a50f-2bfbefb5ddbb.

"Can You Comment on the Meaning of These Moments of Laughter and Applause?" 2025. Gemini. April 4, 2025. https://gemini.google.com.

"Can You Watch This Video and Analyse the Audience Reactions?". 2025. Gemini. April 4, 2025. https://gemini.google.com.

Chun, Wendy Hui Kyong. 2017. "Digital and New Media." In *A Companion to Critical and Cultural Theory*, edited by Imre Szeman, Sarah Blacker and Justin Sully. Hoboken, NJ: Wiley Blackwell.
Chun, Wendy Hui Kyong. 2018. "Queerying Homophily Muster Der Netzwerkanalyse." *Zeitschrift Für Medienwissenschaft* 10 (18-1): 131–148. https://doi.org/10.14361/zfmw-2018-0112.
Chun, Wendy Hui Kyong. 2021. *Discriminating Data: Correlation, Neighborhoods, and the New Politics of Recognition*. Cambridge, Massachusetts: MIT press.
CTV Your Morning. 2024. "Anesti Danelis' New Show 'Artificially Intelligent." https://www.youtube.com/watch?v=K6-lrzu2Qu0.
Crary, Jonathan. 2005. *Techniques of the Observer: On Vision and Modernity in the Nineteenth Century*. Nachdr. October Books. Cambridge, Mass.: MIT Press.
Double, Oliver. 2014. *Getting the Joke: The Inner Workings of Stand-up Comedy*. London: Bloomsbury Publishing Plc. https://doi.org/10.5040/9781408177686.
Fei, Hao, Yuan Yao, Zhuosheng Zhang, Fuxiao Liu, Ao Zhang, and Tat-Seng Chua. 2024. "From Multimodal LLM to Human-Level AI: Modality, Instruction, Reasoning, Efficiency and beyond." Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024) Torino, Italia, May.
Fitzpatrick, Dan. 2024. "How ChatGPT Voice Has Made The World More Accessible." *Forbes*, 5 Oct https://www.forbes.com/sites/danfitzpatrick/2024/10/05/how-chatgpt-voice-has-made-the-world-more-accessible/.
Folger Shakespeare Library. 2021. "How We Hear Shakespeare's Plays, with Carla Della Gatta." Shakespeare Unlimited Podcast. July 20, 2021. https://www.folger.edu/podcasts/shakespeare-unlimited/hearing-a-play-della-gatta/.
Gadsby, Hannah. 2018. Nanette. Olb, Jon, and Madeleine Perry, dirs. Netflix. www.netflix.com.
Gemeinboeck, Petra. 2021. "The Aesthetics of Encounter: A Relational-Performative Design Approach to Human-Robot Interaction." *Frontiers in Robotics and AI* 7: 577900. https://doi.org/10.3389/frobt.2020.577900.
Gemeinboeck, Petra, and Rob Saunders. 2023. "Hybrid Entanglements: A Posthuman Dramaturgy for Human-Robot Relationships." *ISEA 2023: Symbiosis* Paris.
Giaccardi, Elisa, and Johan Redström. 2020. "Technology and More-Than-Human Design." *Design Issues* 36 (4): 33–44. https://doi.org/10.1162/desi_a_00612.
Giuffre, Liz. 2021. "From Nanette to Nanettflix – Hannah Gadsby's Challenge to Existing Comedy Convention." *Comedy Studies* 12 (1): 29–39. https://doi.org/10.1080/2040610X.2020.1850102.
Goltz, Dustin Bradley. 2017. *Comic Performativities: Identity, Internet Outrage, and the Aesthetics of Communication*. 1st ed. London and New York: Routledge. https://doi.org/10.4324/9781315181967.
Gray, Carson, Paris Myers, and Naomi T. Fitter. 2020. "Read the Room, Robot! Exploring Audiovisual Methods to Improve the Effectiveness of Robotic Comedians." In: Proceedings of the IROS Workshop on Social AI for Human-Robot Interaction of Human-Care Service Robots.
Heim, Caroline. 2016. *Audience as Performer: The Changing Role of Theatre Audiences in the Twenty-First Century*. London and New York: Routledge.
Jenkins, Henry. 1992. *What Made Pistachio Nuts?: Early Sound Comedy and the Vaudeville Aesthetic. Film and Culture*. New York: Columbia University Press.
Kahekashan. 2024. "*ChatGPT Now Supports Video Inputs: All Details*." December 14. https://www.thehansindia.com/tech/chatgpt-now-supports-video-inputs-all-details-929713.
Khoo, Lin Sze., Mei Kuan Lim, Chun Yong Chong, and Roisin McNaney. 2024. "Machine Learning for Multimodal Mental Health Detection: A Systematic Review of Passive Sensing Approaches." *Sensors* 24 (2): 348. https://doi.org/10.3390/s24020348.
Krefting, Beck. 2024. "Hannah Gadsby: Emotional Capital and Affective Economies in Stand-Up Comedy." *European Journal of American Studies* 19 (3). https://doi.org/10.4000/12avl.
Lawton, Megan. 2024. "AI and Stand up Comedy." *BBC*, August 15. https://www.bbc.com/news/articles/c1jll2k0gewo.
Lee, Young Suk 2025. ""Kissing Birds": Mediating Human-Human Interaction Through Poetic Engagements with Technology." In *Social Robotics. ICSR + AI 2024. Lecture Notes in Computer Science*, edited by Palinko, O, et al., 15561. Singapore: Springer. https://doi.org/10.1007/978-981-96-3522-1_5.

Lintott, Sheila. 2020. “Stand-Up Comedy and Mental Health: Critiquing the Troubled Stand-Up Stereotype.” In *The Dark Side of Stand-Up Comedy*, edited by Patrice A. Oppliger and Eric Shouse, 197–222. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-37214-9.

Lucie, Sarah. 2021. “Human Objecthood in the Datasphere: The Body as Resistant Material in Twenty-First-Century Media.” *Theatre Journal* 73 (3): 319–337. https://doi.org/10.1353/tj.2021.0070.

McLuhan, Marshall. 1964. *Understanding Media: The Extensions of Man*. New York: New American Library.

Meany, Michael. 2016. "Comedy, Creativity, Agency: The Hybrid Individual." In *The Creative System in Action: Understanding Cultural Production and Practice*, edited by Phillip McIntyre, Janet Fulton and Elizabeth Paton, 169–184. London: Palgrave Macmillan UK.

Metzger, Sean. 2021. “Editorial Comment: On the Possibilities of AI, Performance Studies, and Digital Cohabitation.” *Theatre Journal* 73 (3): xiii–xixx. https://doi.org/10.1353/tj.2021.0065.

Milner Davis, Jessica. 2014. “Bergson and the Theory of the Comic.” In *Encyclopedia of Humor Studies*, edited by Salvatore Attardo, 1, 77–80. Thousand Oaks CA: Sage. https://us.sagepub.com/en-us/nam/encyclopedia-of-humor-studies/book235990.

Mintz, Lawrence E. 1985. “Standup Comedy as Social and Cultural Mediation.” *American Quarterly* 37 (1): 71–80. https://doi.org/10.2307/2712763.

Mirowski, Piotr Wojciech, Juliette Love, Kory W. Mathewson, and Shakir Mohamed. 2024. “A Robot Walks into a Bar: Can Language Models Serve as Creativity Support Tools for Comedy? An Evaluation of LLMs’ Humour Alignment with Comedians.” May. http://arxiv.org/abs/2405.20956.

Mirowski, Piotr Wojciech, Boyd Branch, and Kory Wallace Mathewson. 2025. “The Theater Stage as Laboratory: Review of Real-Time Comedy LLM Systems for Live Performance.” 1st Workshop on Computational Humor (CHum).

Netflix Is A Joke. 2021. “Hannah Gadsby on Being a Woman in a Man’s World.” https://www.youtube.com/watch?v=IuoTAkn-0DQ.

Oxford English Dictionary. 2025. “Audience (*n.*).” March. https://doi.org/10.1093/OED/5022349431.

Oxford English Dictionary. 2025. “Auditorium (*n.*).” March. https://doi.org/10.1093/OED/5023695169.

Quirk, Sophie. 2015. *Why Stand-up Matters: How Comedians Manipulate and Influence*. 1st ed. London: Bloomsbury Publishing Plc. https://doi.org/10.5040/9781472578969.

Reason, Matthew, Lynne Conner, Katya Johanson, and Ben Walmsley, eds. 2022. “Routledge Companion to Audiences and the Performing Arts.” In *Audience Research Ser*. Milton: Taylor & Francis Group.

Rozendaal, Marco C., Jered Vroon, and Maaike Bleeker. 2025. “Enacting Human–Robot Encounters with Theater Professionals on a Mixed Reality Stage.” *ACM Transactions on Human-Robot Interaction* 14 (1): 1–25. https://doi.org/10.1145/3678186.

Swaminathan, Janani Jane., Akintoye, Marlena R. Fraune, and Heather Knight. 2021. 2021. “Robots That Run Their Own Human Experiments: Exploring Relational Humor with Multi-Robot Comedy.” 2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN), August 8–12.

Tang, Xiaoyu, Jiazheng Huang, Yixin Lin, Ting Dang, and Jintao Cheng. 2025. “Speech Emotion Recognition via CNN-Transformer and Multidimensional Attention Mechanism.” *Speech Communication* 171 (June): 103242. https://doi.org/10.1016/j.specom.2025.103242.

Turing, Alan Mathison. 1950. “Computing Machinery and Intelligence.” *Mind* LIX (236): 433–460. https://doi.org/10.1093/mind/LIX.236.433.

Ullrich, Daniel, Andreas Butz, and Sarah Diefenbach. 2020. “The Eternal Robot: Anchoring Effects in Humans’ Mental Models of Robots and Their Self.” *Frontiers in Robotics and AI* 7: 546724. https://doi.org/10.3389/frobt.2020.546724.

Walton, Robert Ellis. 2021. “Theatres of Artificial Intelligence and the Overlooked Performances of Computing.” *Theatre Journal* 73 (3): 279–298. https://doi.org/10.1353/tj.2021.0067.

“What Is Gemini and How It Works.” n.d. Gemini. Accessed April 6, 2025. https://gemini.google/overview/?hl=en.

"What Can You Tell Me about Nanette's Audience?" 2025. *ChatGPT*. February 14. https://chatgpt.com/g/g-TDjzPM7vk-video-summary-and-analysis/c/67aec5e1-ee1c-8012-a50f-2bfbefb5ddbb.
Winters, Thomas. 2021. "Computers Learning Humor is no Joke." *Harvard Data Science Review* 3 (2). https://doi.org/10.1162/99608f92.f13a2337.
Walker, Julia A. 2022. *Performance and Modernity: Enacting Change on the Globalizing Stage*. 1st ed. Cambridge: Cambridge University Press. https://doi.org/10.1017/9781108966870.