

Roddy Cowie*

Computational research and the case for taking humor seriously

<https://doi.org/10.1515/humor-2023-0021>

Received January 28, 2020; accepted February 5, 2023; published online March 17, 2023

Abstract: Computational research underscores the complex abilities underlying humor. Two decades of work have achieved substantial progress in some areas, notably systems that make jokes; detecting and generating laughter; and using irony in interactions. Sophisticated evaluations clarify both strengths and limitations. The achievements illuminate specific abilities, but also expose unsolved problems. The way humor pervades life is harder to match than self-contained episodes. Learning techniques are powerful, but providing the data they need is daunting. The medium is no longer simply verbal, but other modalities present deep challenges, such seeing the humor in a situation. There are real applications, but the most striking still depend on human support. Achievements and limitations together underscore the scale of the challenges involved in understanding humor.

Keywords: artificial intelligence; humor; pervasive humor

1 Introduction

Research fields can often make different kinds of contribution when they are viewed from different perspectives. This paper views computational research on humor from the perspective of an outsider, interested in computational research primarily for the light it sheds on human abilities.

From that perspective, one of the contributions that computational research often makes is to foster appreciation of abilities that are liable to be undervalued. The tendency to disrespect humor is well known. Traditions going back to Plato portray it as an unedifying part of life, intellectually lazy and morally dubious (see, e.g., Morreall 1997). There are various established ways of making the case for a much higher evaluation (see, e.g., Morreall 2020). This paper tries to indicate that computational research makes a distinctive addition to that case.

*Corresponding author: Roddy Cowie, Queen's University, Belfast, Northern Ireland,
E-mail: r.cowie@qub.ac.uk. <https://orcid.org/0000-0003-3480-2223>

 Open Access. © 2023 the author(s), published by De Gruyter.  This work is licensed under the Creative Commons Attribution 4.0 International License.

There is a substantial research effort in the field, and there is no question of reflecting the whole of it in a short article. However, a selective sketch may give a sense of its contribution to the case for taking humor seriously. It has taken sustained and impressive research to match a subset of the abilities involved in everyday humorous interactions. Knowing that vividly underlines the breadth and flexibility of the way humans routinely use humor.

A striking part of the case involves comparison with work on matching other human abilities – artificial intelligence in a broad sense. It has studied abilities that people find impressive when humans display them – proving theorems in logic, or playing master-level chess – and created machines that outperform most humans (or all). But although people rarely set much store by their ability to defuse a tricky interaction with a touch of humor, that turns out to be enormously more difficult to match. The comparison raises questions about the way we evaluate humor relative to other abilities. Similar things can be said about other areas, such as perception and emotion: the point here is that they can be said about humor, too.

The scale of the challenge was not unforeseen. Concerted research in the area took shape around the turn of the century. A useful marker is the workshop on verbal humor convened by Hulstijn and Nijholt (1996). A broader picture quickly emerged. It is well reflected in the April Fools' Day Workshop on Computational Humour, in 2002. It involved significant figures from several fields – Attardo from linguistics, Hofstadter from cognitive science, Ortony from emotion research, Ruch from the psychology of humor. At that stage Oliviero Stock, the host, gave a summary that stands the test of time: "We are convinced that deep modeling of humor in all of its facets is not something for the near future. The phenomena are too complex; humor is one of the most sophisticated forms of human intelligence." (Stock and Strapparava 2002).

The review in this paper tries to indicate that, on the basis of research since, Stock is well vindicated. Other summaries take different perspectives, with prominent contributions from the workshop co-organizer, Anton Nijholt (e.g. Nijholt et al. 2017). However, the research is well worth considering from varied angles.

1.1 Groundwork for computational research

The April Fools' Day Workshop conceptualized computational research on humor as a field that would naturally engage with several others.

HCI gave research practical motives. In keeping with the case for taking humor seriously, it provided evidence against the assumption that introducing humor into serious tasks could only deflect attention, and impair performance (Morkes et al.

1999). Raskin (2002) summarized the applications that were foreseen, using a list due to Stock, with two additions of his own:

- business world applications (such as advertising, e-commerce, etc ...);
- general computer-mediated communication and human-computer interaction;
- increase the friendliness of natural language interfaces;
- edutainment and autonomous agents systems;
- customer acceptance enhancement; and
- humor detection.

Linguistics provided a highly developed theoretical analysis of patterns that generate amusement, derived from the General Theory of Verbal Humour of Attardo and Raskin (1991), and represented in the workshop by Attardo (2002) and Raskin (2002). Broadly, the patterns involve partial parallels between two semantic structures. The starting point for a joke is a reasonably natural way of putting most of the elements on one side into correspondence with elements on the other. However, the correspondence invites us to pair up the remaining elements on either side: and that pairing is strikingly anomalous. The combination of match and anomaly is what prompts amusement. Predictions from that framework were supported by empirical research (Ruch et al. 1993). Conference papers reflected work which had already used related ideas to generate riddles, the *Jape* system (Binsted and Ritchie 1997). Related ideas were applied to amusing visual stimuli (Giora 2002).

Empirical contributions highlighted the significance of context. Ruch (2002) stressed that different kinds of humor characterize different kinds of personality, and are received differently by them. Ben Ze'ev (2002) considered the different roles that humor plays in different media, and the enhanced role it plays in online communication.

Research on emotion also featured, though less systematically. The view that amusement is a form of pleasure was occasionally explicit (Giora 2002), but often implicit. Nijholt (2002) proposed a more theoretical conception: humor as a form of appraisal, involving patterns that are anomalous, but not harmful. He also noted its role in expressing emotion acceptably: “if [artificial agents] are going to show emotion, we surely hope that they would show a little humour too.”

The range of resources that were considered underlines the scale of the challenge. Looking at the result, again broadly, it seems fair divide it into two. Particular directions have developed strongly. The achievements there serve as landmarks. Elsewhere, the significant achievement is clarifying targets that are important, but challenging. The next two sections look at those in turn. That provides at least a broad basis for assessing how attempts to match these particular human abilities contribute to understanding ourselves.

2 Landmark developments

By the nature of humor, it would be surprising if research divided neatly. Nevertheless, it seems fair to pick out four areas where research can claim substantial achievements.

2.1 Comedians

‘Comedian’ is a natural term for someone whose goal – at a given time – is to make people amused. Most people do not adopt the role very often: humor usually arises in the course of pursuing other kinds of goal (or abandoning them). In contrast, it has been very prominent in computational research.

The field’s early successes were in that vein. For example, the *Jape* system (Ritchie 2003) produced punning riddles such as:

What is the difference between a pretty glove and a silent cat?

One is a cute mitten, the other is a mute kitten.

The evaluation was revealing: *Jape*’s best outputs were compared with jokes in children’s books, and children found them about equally amusing (Binsted et al. 1997).

Varied formal patterns were explored. The HAHAcronym generator (Stock and Strapparava 2002) generated amusing misinterpretations of acronyms by choosing words with associations diametrically opposed to the originals: so, MIT becomes “Mythical Institute of Theology.” In similar vein, Valitutti (2009) explored creative variation of familiar expressions: so “Tomorrow is Another Day” became “Tomorrow is Another Bay.” Tinholt and Nijholt (2007) made a significant innovation: the system interacted with humans, and took up opportunities to make formulaic jokes that arose from what they said. Its formula detected the referential ambiguity in a sentence like “The cops arrested the demonstrators because they were violent,” and highlighted the anomalous reading: “What? The cops were violent?”

Several of those systems have a second level. The first generates texts in a form that can be funny: the second evaluates its outputs, using various criteria. HAHAcronym checks for contrast in meaning between the original phrase and its alternative; Hempelmann’s (2003) punning system checks how close the punning sound is to the original; and rules involving the familiarity, concreteness and

imageability of the descriptions were used to select outputs from *Jape* that were likely to amuse raters. It seems possible that humans do something similar.

Jape's output was compared with material from children's joke books: matching comic performance is a natural follow-on. Efforts in that direction are summarized by Nijholt (2018). They are strikingly unlike earlier comedic work in some respects, strikingly like it in others.

Nijholt notes a striking dissimilarity: the performance systems make limited use of techniques for generating jokes. Scripts are generally provided by humans; the system's task is delivery. The striking similarity is that there are staple patterns to guide delivery as well as content.

Manzai comedy provides one set of staples, involving exaggerated behaviors, unexpected actions, interactions involving imitation or misunderstanding, and some kinds of verbal humor. Examples involve interactions between a robot and humans (Kishi et al. 2014), and between two robots (Umetani et al. 2017). Another robot comedian used four 'performative gestures' derived from the 'stand-up' style: an opening welcome gesture; a reprise ("I said hello"); pointing to the audience; and applause elicitation at the end of the performance (Katevas et al. 2015). Underpinning all of those are various levels of interaction with the audience: looking at individuals, detecting amusement, using their prompts to choose among available topics.

The state of the art in this area is a revealing marker. We might naively assume that since a General Theory of Verbal Humour was articulated several decades ago, computational research would by now have translated it into scintillating humorists. The actual state of the art underlines Stock's assessment: "deep modeling of humor ... is not something for the near future."

2.2 Laughter

Laughter is a natural step beyond textual humor. It is the obvious sign of amusement, and effective performance comedy depends on detecting it. Since it is infectious, it is also a means of inducing amusement. But as with comedy, exploring the subject led well beyond the obvious.

Basic parts of the work could, in principle, have been done for pure scientific curiosity; but in practice, computational aims drove it. On one side of that is machine perception: finding ways to derive useful descriptions from sensors. That depends on data, and so computational needs prompted the creation of laughter databases (see Dupont et al. 2016). On the other side, there is the problem of working out when and how an agent should laugh.

Recognizing laughter sounds simple, but it turned out to be intricate (Kantharaju et al. 2019). A first stage is separating standard speech from pure laughter, voiced and unvoiced; and speech overlaid with laughter. Out of a range of tools, visual as well as acoustic, the most effective there was a battery of acoustic features used in other emotion detection contexts, eGeMAPS. For the next step, identifying the emotion carried, only voiced laughter and laughter in speech gave useful indications. They came from combining eGeMAPS features with visual information. That combination also distinguished between spontaneous and acted laughter.

Moving to another level again, humans judge that laughs differ in intensity. That judgment turns out to depend on multiple acoustic features, including voicing; but also duration, the existence of harmonic structure, and different features related to pitch (Rychlowska et al. 2019). Intensity then features in another level (Curran et al. 2018). High-intensity laughs appear to have an intrinsic meaning; but low-intensity laughs acquire meaning from the context (so their perceived meaning will change if they are inserted into a new context).

The inverse problem is to establish how and when to generate laughter. Recent studies point to three natural ideas. Mancini et al. (2017) describe an agent that laughs in response to humor in external stimuli (pieces of music). Their question is how the laughter affects the perception of a human who is also listening. The agent in Niewiadomski et al. (2013) uses a contrasting strategy, qualified imitation: it laughs in response to its human companion's laughs, somewhat (but not exactly) similarly. Piot et al. (2014) take a third approach: their agent's laughter is governed by a learning algorithm, which takes account of situational variables and a human partner's responses. It seems reasonable to assume that all three strategies capture something significant: combining them is another level of complexity.

Laughter illustrates the general theme particularly sharply. It is an aspect of life that is prone to be regarded as crude, even animalistic. But trying to match what we do reveals that it is bewilderingly complex. It only appears simple because the subtleties are handled outside conscious awareness.

2.3 Irony

Irony came to notice as a feature of online texting. Work there (e.g. Valitutti and Veale 2015) showed how simple devices – scare quotes and dramatic adjectives – could induce a ‘valence shift’ in attitudes to an underlying message (e.g. The vegetables are mixed in “healthful” salads). The topic attracted considerable interest in human-robot interaction. It has several attractions there. It is a form that can be integrated into communication, rather than standing alone, as a joke does. It also relates to a particular problem, which is orienting a human to a robot's mix of

abilities and limitations. A ‘valence shift’ that makes a joke of failures applies neatly to that problem.

A dramatic demonstration was provided by two robots, which responded to failures in contrasting ways (Mirnig et al. 2016). Both were humorous, but in one it took the form of *schadenfreude*, mocking the other’s failures, where the other used irony to mock its own. The second was rated more likeable. That is a useful counter to the assumption in early research that humor *per se* is the key to evoking positive reactions. Clearly, the kind of humor matters too.

‘Irony man’ (Ritschel et al. 2019) used prior theory to develop the approach. The basic tool is ideational reversal, where “the intended meaning arises as a result of negation of a chosen element of the literally expressed meaning or the pragmatic import of the entire utterance.” Instead of ‘scare quotes’, the key element is marked by signals in various modalities – exaggerated or understated language, prosody, gaze, facial expression, and gesture.

A further feature is that the robot uses irony in the context of a discussion with a person and chooses when to use it. A simple principle governs the choice. Irony is used to signal agreement with an emotive statement made by the user, by repeating the statement with an inversion that is signaled as ironic. For instance, the system responds to “I hate raining” by saying “Super! I utterly *love* raining” – with suitable markers on ‘love,’ The simplicity is valuable, because it allows the kind of learning that has already been mentioned in connection with laughter, aimed at recognizing contexts where the tactic is likely to work.

An unusually broad evaluation indicates that the basic form does work. It compared versions of the robot with and without irony. The ironic version was perceived as more natural, and more pleasant on a measure covering a range of subjective evaluations – more stylish, presentable, attractive, appealing, and likeable.

2.4 Outcomes

Irony research illustrates a last major area of progress: clarifying what systems achieve. It is useful to draw a rough line between two generations of work.

In the first generation, the issue was global: whether computational humor could have positive effects. Varied approaches converged on a positive conclusion. Babu et al. (2006) studied responses to a virtual receptionist, Marve. Users could choose how, and how much, they interacted with him. Listening to his jokes was much the most popular option, accounting for 50% of interaction time. Dybala et al. (2009) compared two versions of a conversational system, identical except that one (Pundalin) told a pun every third conversational turn. Pundalin received

predominantly positive affective responses, the humorless system predominantly negative. Similar evaluations continue: e.g., Mohammadi et al. (2019) compared agents that play dice with users, one which interacts humorously, the other competitively. The humorous agent was judged more animated, likeable, and empathic, and more able to experience a range of emotions.

Computational humor certainly can have positive effects. But, notoriously, humor out of place can go badly wrong. The second generation aimed to disentangle appropriate and inappropriate contexts.

Khooshabeh et al. (2011) considered humorous and humorless versions of an agent designed to persuade users. Different users responded differently. When users did not like the humor, it made them less likely to be persuaded. Karpinska-Krakowiak et al. (2014, 2018) highlighted the growing use of pranks in online advertising. Those effects too are unstable: they put off people who identify with the victims of the pranks. With *schadenfreude* humor (Mirnig et al. 2017), the reactions divided along gender lines. Torre et al. (2019) considered a subtler division, using a game where an avatar suggested solutions to a problem, accompanied by a smile. A smile signaled only visually made its suggestions more likely to be accepted, but adding smiling voice made them less so. A degree of neutrality increases trust.

AI offers an obvious response to the finding that both positive and negative outcomes are possible. It is to use machine learning techniques to ensure that humor is properly placed. Examples of that approach have been mentioned, involving laughter and particularly irony (Weber et al. 2018). However, that exposes another level of challenge: how to understand exactly which outcomes are desirable, or how they are signaled; or along which dimensions situations, and users, need to be distinguished. Those point back to the questions about the context of humor that were raised in the April Fool's Day workshop.

3 Emerging targets

Alongside the landmarks are topics that are clearly signposted, but further from satisfying solution. That is rarely because people in the field are unaware of the issues; usually, it is because the challenges are deeper than first impressions suggest.

3.1 Pervasive and emergent humor

Emotion research highlights an issue that extends to humor. The word 'emotion' tends to call to mind distinctive states where emotion interrupts rational thought: 'emergent emotions.' Those contrast with another distinctive state, being

unemotional. However, most of everyday life is neither. The norm is ‘pervasive emotion’, which colors the way people think, and act, and interact, without suppressing reason (Cowie 2010).

Pervasive humor is an implicit challenge reflected in several strands of computational research. Irony is a form of humor that can pervade interaction (unlike constant set piece jokes). Acceptable pervasiveness depends on taking up opportunities that interaction and context offer, either to laugh (Piot et al. 2014) or to make jokes (Wong and Gonzalez 2016). Linking humor to personality recognizes one of the major inputs from psychology at the April Fool’s Day workshop: a pervasive style of humor is one of the things that distinguishes a person (Ruch 2002).

All of those represent progress into the ground between humorless systems and comedians. The limits of progress make a very enlightening point. The middle ground is not a routine compromise between the extremes. Mastering the relevant blends and contingencies is extraordinarily difficult.

3.2 Data

AI in general began with a hypothetico-deductive model: machines derived the implications of principles formulated by humans. That wave – ‘good old-fashioned AI’ – was superseded by techniques based on learning (Haugeland 1989). The general pattern extends to computational humor: hypothetico-deductive approaches were prominent in the April Fool’s Day Workshop, and learning has come increasingly to the fore.

Learning techniques invite a shift from symbolic representations of meaning to probabilistic networks, and that has quite subtle consequences (Hempelmann and Petrenko 2015). More straightforwardly, learning depends on access to large quantities of reliable data. Acquiring it is a much bigger part of the enterprise than people outside the field commonly recognized, not least because the raw material must be labelled to reflect people’s responses to it.

There are substantial bodies of data in some subareas. Online sources have been used to build databases of jokes and funny dialogues, spanning different languages. For example, the HAHA challenge compares systems’ ability to learn from a database of 30,000 annotated Spanish tweets (Chiruzzo et al. 2019); and Blinov et al. (2019) describe a Russian dataset including more than 300,000 short texts. For laughter, the ILHAIRE project left a legacy of annotated databases, spanning different contexts and cultures (Dupont et al. 2016). That effort continues (Ginzburg and Pelachaud 2018).

Other areas use varied forms of data. Radev et al. (2015) considered cartoons and their captions using over 400 cartoons and more than 2 million proposed captions, acquired from *New Yorker* competitions. Oliveira et al. (2016) generated amusing

picture/text combinations using a stock of preselected images and subjects drawn from newspaper headlines. Systems that learn by interaction, mentioned earlier, effectively use another form of data: each item pairs an effort at humor with an evaluation of listeners' responses.

The issue is an illuminating one. Despite the effort involved, available resources only partially address specific subtasks. The degrees of freedom involved in modelling pervasive humor would be of a different order of magnitude. It is hard to imagine the amount of data that would have to be collected, or how it could be labelled. Progress in that direction matters for understanding ourselves, though. Learning involves finding solutions to the problems that a database defines. Restricted databases may well allow solutions quite unlike those that let people find humor in hugely diverse situations. Building databases that direct learning towards human-like solutions is a non-trivial part of the challenge.

3.3 Extending modalities

Theory addressed visual jokes long ago (e.g. Giora 2002), but early computational work was entirely verbal. Extension to other modalities is an ongoing challenge.

Humor generation provides examples that have already been mentioned. In Manzai comedy, posture and gesture function like jokes. In irony, visual and prosodic markers enable verbal humor. The modalities also work together in picture/text combinations. A relatively recent point is that verbal and non-verbal signals are often interchangeable rather than additive (Mirnig et al. 2017). Humor in video games uses multiple modalities, up to and including visual slapstick (Hessler 2019).

Multimodal perception is much harder. It is used to detect a human's amusement, or laughter. Recognizing what is amusing is a different matter. The study of New Yorker cartoons is illuminating. The cartoons only entered the analysis by way of verbal descriptions identifying the objects in them (which were provided by the researchers). That is far short of implementing theories of humor in cartoons, like Giora's; or registering the slapstick, or funny incongruity, in video clips.

Again, that relates to pervasive humor. One of the everyday benefits of a sense of humor is the ability to see the funny side of a situation that could be distressing. Matching that would need major advances in perception.

The issue is not new. Nijholt's paper in the April Fool's Day workshop considered humor as an appraisal. Its business is detecting that a situation violates expectations embedded in plans, or standard images of things: but it is not threatening. The associated action tendency could plausibly involve suppressing responses to the apparent threat, and giving positive signals. That portrays humor as critiquing our

representations and projections. It implies a kind of double perception, of reality and of the way we represent it. Implementing that is a long-term project.

3.4 Applications

Applicability was crucial for the April Fool's Day conference, and it is intimately linked to the case for taking humor seriously. Applications have materialized, though not quite as anticipated.

Stock's list did not mention gaming, but it is an application that became established early on (Dormann and Biddle 2009; Grönroos 2013). Applications in education, which did feature, followed a related path: humor is a regular element of games designed to teach (e.g., Pezzullo et al. 2017). The usefulness of the technology has to be assessed case by case (Foldager et al. 2016). Another application area involves 'artificial companions', particularly to supplement human support in long term care environments. In a forum on the subject, Romano (2007) began her contribution: "A truly good friend is the one that makes you laugh." It has been seen as a role where agents could use techniques derived from stand-up comedy (e.g. Kishi et al. 2014).

However, the outstanding applications reflect Stock's concern to "increase the friendliness of natural language interfaces." Chatbots have become an established form of interface (Grudin and Jacques 2019), and humor is an established tool. Significantly, it is a feature of the major IT companies' chatbots, Cortana, Siri, and Alexa.

The major chatbots cast a revealing light on the state of the art. On one side, their use of humor puts it beyond doubt that the technology has applications. On the other hand, it highlights the limits of the technology. The outstanding examples rest on the same theme. Humor is part of larger wholes; and assuming that it can be detached from the wholes, except in special circumstances, is misunderstanding it.

In Luger and Sellen's memorable description, using a chatbot can be "like having a really bad PA" (2016). They found that systems could not live up to the expectations they raised, resulting in disillusionment and curtailed interaction. Humor played a part in that. It encouraged playful interactions early on. A detailed study by Liao et al. (2018) noted a similar pattern of initial playfulness, which humor promoted. However, the initial interactions did not simply generate positive affect. Users understood them as a sign of a larger whole, involving a range of competences: and reacted negatively when they realized that the system did not have those qualities.

Alongside the work on software agents, there is emerging work on humor in interactions with physical robots (including work on irony that was mentioned

earlier). It raises an intriguing question: whether different embodiments require different humors.

Practical experience also underlines the importance of another whole: that is, personality. The humor in the major companies' chatbots is not wholly computational. The systems that deal directly with human users accumulate large amounts of data and learn from it. But they are backed up by human teams, who monitor and develop response patterns. The process, described by Elgan (2016), is "a lot like novel writing. There's no story or plot, but there is one character, which has a gender, an age, an educational level, a particular sense of humor and more. That human-created character or personality has to remain consistent across queries and contexts."

The linkage between personality and humor is no new discovery: Ruch (2002) underlined it at the April Fool's Day workshop. What has emerged is how challenging it is to ensure that a sense of humor is internally coherent and coheres with a wider personality. For that, the best-resourced teams still rely on creative human intuition.

4 Conclusion

There are other ways of conceptualizing AI, but this paper has tried to reflect one that seems useful. In humor as elsewhere, it takes descriptions of people from other sources, and tries to translate them into systems that match their achievements. The process exposes the limits of the descriptions. A major part of the reason is that faced with a system doing things we expect humans to do, we tend to notice what is missing. With descriptions, we make allowances: we expect them to be partial. However, a functioning system invites comparison with a functioning human being. That can sharpen our understanding of the descriptions: it reveals that the things they say are the tip of an iceberg, most of which consists of things have been left unsaid.

From that perspective, AI offers traditional approaches a powerful ally. It has priorities that need to be reckoned with. It must be opportunistic and tends to take up descriptions that lend themselves to implementation. Having its roots in the 'mechanization of thought processes', it tends to favor thought-like models, at least initially. Its other staple, learning, directs it towards areas where data can be accumulated. To pay its way, it must prioritize developments with potential applications. Those pressures need to be factored into understanding of what it does. But once they are, AI provides a striking perspective on the strands where development can take place – showing what is unexpectedly simple, and what is far more challenging than it seemed. Most strikingly, 'the quest for artificial general intelligence' (Grudin and Jaques 2019) underlines how little we understand human

ability to combine vastly diverse abilities in a coherent way. We can simulate some very interesting parts: it is the whole that eludes us.

Computational humor offers an intriguing example of that pattern. It addresses an aspect of life that has often been considered crude. But evidence from sustained efforts to match it favors an assessment closer to Stock's intuition that "humor is one of the most sophisticated forms of human intelligence." It also indicates that, as elsewhere, it is the whole that eludes us. Metaphorically, artificial agents that attempt humor are a useful companion to Socrates: they bring us face to face with the limitations of what we think we know about ourselves.

References

- Attardo, Salvatore. 2002. Formalizing humor theory. In *Proceedings of Twente Workshop on Language Technology*, vol. 20, 1–10. Enschede: University of Twente.
- Attardo, Salvatore & Victor Raskin. 1991. Script theory revis(it)ed: Joke similarity and joke representation model. *Humor* 4(3–4). 293–347.
- Ben-Ze'ev, Aaron. 2002. Humor in online romantic relationships. In *Proceedings of Twente Workshop on Language Technology*, vol. 20, 71–76. Enschede: University of Twente.
- Babu, Sabarish, Stephen Schmutge, Tiffany Barnes & Larry Hodges. 2006. What would you like to talk about? An evaluation of social conversations with a virtual receptionist. In Jonathan Gratch, Michael Young, Ruth Aylett, Daniel Ballin & Patrick Olivier (eds.), *IIVA 2006*, 169–180. Berlin: Springer.
- Binsted, Kim & Graeme Ritchie. 1997. Computational rules for generating punning riddles. *Humor* 10(1). 25–76.
- Binsted, Kim, Helen Pain & Graeme Ritchie. 1997. Children's evaluation of computer-generated punning riddles. *Pragmatics and Cognition* 5(2). 309–358.
- Blinov, Vladislav, Valeria Bolotova-Baranova & Pavel Braslavski. 2019. Large dataset and language model fun-tuning for humor recognition. In *Proceedings of 57th Annual Meeting of the Association for Computational Linguistics*, 4027–4032.
- Chiruzzo, Luis, Santiago Castro, Mathias Etcheverry, Diego Garat, Juan José Prada, and Aiala Rosá. 2019. Overview of HAHa at IberLEF 2019: Humor analysis based on human annotation. In *IberLEF@SEPLN, Bilbao*. 132–144.
- Cowie, Roddy. 2010. Describing the forms of emotional colouring that pervade everyday life. In P. Goldie (ed.), *The Oxford handbook of philosophy of emotion*, 63–94. Oxford: Oxford University Press.
- Curran, William, Gary McKeown, Magdalena Rychlowska, Elisabeth André, Johannes Wagner & Florian Lingenfelser. 2018. Social context disambiguates the interpretation of laughter. *Frontiers in Psychology* 8. 2342.
- Dormann, Claire & Robert Biddle. 2009. A review of humor for computer games: Play, laugh and more. *Simulation & Gaming* 40(6). 802–824.
- Dupont, Stéphane, Hüseyin Çakmak, Will Curran, Thierry Dutoit, Jennifer Hofmann, Gary McKeown, Pietquin Olivier, Tracey Platt, Willibald Ruch & Jérôme Urbain. 2016. Laughter research: A review of the ILHAIRE project. In A. Esposito & L. C. Chain (eds.), *Toward robotic socially believable behaving systems*, vol. 1, 147–181. New York: Springer.

- Dybala, Pawel, Michal Ptaszynski, Rafal Rzepka & Kenji Araki. 2009. Humoroids: Conversational agents that induce positive emotions with humor. *AAMAS'09 Proceedings of 8th International Conference on Autonomous Agents and Multiagent Systems*, 2, 1171–1172. Budapest: ACM.
- Elgan, Mike. 2016. Virtual assistants: They're made out of people! Why Siri, Cortana and Alexa are only semi-fictional characters. *Computerworld*, 1 February 2016 Retrieved from <https://www.computerworld.com/article/3028003/virtual-assistants-theyre-made-out-of-people.html>.
- Foldager, Nicolai, Hans Hansen, Mikkel Skovsmose Tewes & Thomas Bjørner. 2016. Designing an engaging and informative application about first aid: Gamification and humor as design elements in a serious game. In *International Conference on Smart Objects and Technologies for Social Good*, vol. 78–87. Cham, Switzerland: Springer.
- Giora, Rachel. 2002. Optimal innovation and pleasure. In *Proceedings of Twente Workshop on Language Technology*, vol. 20, 11–28. Enschede: University of Twente.
- Grönroos, Anne-Marie. 2013. *Humour in video games: Play, comedy, and mischief*. Master's Thesis, Media Lab Helsinki.
- Grudin, Jonathan & Richard Jacques. 2019. Chatbots, humbots, and the quest for artificial general intelligence. In *Proceedings of 2019 CHI Conference on Human Factors in Computing Systems*, 1–11.
- Ginzburg, Jonathan & Catherine Pelachaud (eds.). 2018. *Proceedings of laughter workshop*, vol. 58–63. Paris: Sorbonne Université.
- Haugeland, John. 1989. *Artificial intelligence: The very idea*. Cambridge, Massachusetts: MIT press.
- Hempelmann, Christian. 2003. *Paronomasic puns: Target recoverability towards automatic generation*. Ph.D. Dissertation. Purdue University.
- Hempelmann, Christian & Max Petrenko. 2015. An AI for humorously reframing interaction narratives with human users. In *International Conference on Distributed, Ambient, and Pervasive Interactions*, 651–658. Cham, Switzerland: Springer.
- Hessler, Daniel. 2019. Whose pain is it, anyway? On avatar embodiment, slapstick performances, and virtual pain. *Comedy Studies* 10. <https://doi.org/10.1080/2040610X.2019.1692550>.
- Hulstijn, Joris & Anton Nijholt. 1996. *Automatic interpretation and generation of verbal humor*. Enschede, Universiteit Twente: Faculteit Informatica.
- Karpinska-Krakowiak, Malgorzata & Artur Modlinski. 2014. Prankvertising—pranks as a new form of brand advertising online. *Modern Management Review* 21(3). 31–44.
- Karpinska-Krakowiak, Malgorzata & Artur Modlinski. 2018. The effects of pranks in social media on brands. *Journal of Computer Information Systems* 58(3). 282–290.
- Kantharaju, Reshmashree B., Fabien Ringeval & Laurent Besacier. 2019. Automatic prediction of affective laughter from audiovisual data. In *Proceedings of 20th ACM International Conference on Multimodal interaction*, 220–228.
- Katevas, Kleomenis, Patrick Healey & Matthew Harris. 2015. Robot Comedy Lab: Experimenting with the social dynamics of live performance. *Frontiers in Psychology* 6. Article No. 1253.
- Kishi, Tatsuhiro, Nobutsuna Endo, Takashi Nozawa, Takuya Otani, Sarah Cosentino, Massimiliano Zecca, Kenji Hashimoto & Atsuo Takanish. 2014. Bipedal humanoid robot that makes humans laugh with use of the method of comedy and affects their psychological state actively. *2014 IEEE International Conference on Robotics & Automation*, 1965–1970. New York: IEEE.
- Khooshabeh, Peter, Cade McCall, Sudeep Gandhe, Jonathan Gratch & James Blascovich. 2011. Does it matter if a computer jokes? In *CHI '11 Extended Abstracts on Human Factors in Computing Systems*. 77–86.
- Liao, Q. Vera, Muhammed Mas-ud Hussain, Praveen Chandar, Matthew Davis, Yasaman Khazaeni, Marco Crasso, Dakuo Wang, Michael Muller, Sadat Shami & Werner Geyer. 2018. All work and no

- play? Conversations with a question-and-answer chatbot in the wild. In *Proceedings of CHI 2018*, paper 3, 1–13.
- Luger, Ewa & Abigail Sellen. 2016. “Like having a really bad PA”: The gulf between user expectation and experience of conversational agents. In *Proceedings of CHI 2016*, 5286–5297.
- Mancini, Maurizio, Beatrice Biancardi, Florian Pecune, Giovanna Varni, Yu Ding, Catherine Pelachaud, Gualtiero Volpe & Antonio Camurri. 2017. Implementing and evaluating a laughing virtual character. *ACM Transactions on Internet Technology* 17(1). Article No. 3.
- Mirnig, Nicole, Susanne Stadler, Gerald Stollnberger, Manuel Giuliani & Manfred Tscheligi. 2016. Robot humor: How self-irony and Schadenfreude influence people’s rating of robot likability. In *25th IEEE International Symposium on Robot and Human Interactive Communication*. <https://doi.org/10.1109/ROMAN.2016.7745106>.
- Mirnig, Nicole, Gerald Stollnberger, Manuel Giuliani & Manfred Tscheligi. 2017. Elements of humor: How humans perceive verbal and non-verbal aspects of humorous robot behavior. In *Proceedings of 2017 ACM/IEEE International Conference on Human-Robot Interaction*, 211–212.
- Mohammadi, Hadi Beik, Nikoletta Xirakia, Fares Abawi, Irina Barykina, Krishnan Chandran, Gitanjali Nair, Cuong Nguyen, Daniel Speck, Tayfun Alpay, Sascha Griffiths, Stefan Heinrich, Erik Strahl, Cornelius Weber & Stefan Wermter. 2019. Designing a personality-driven robot for a human-robot interaction scenario. In *2019 International Conference on Robotics and Automation (ICRA)*. 4317–4324.
- Morkes, John, Hadyn K. Kernal & Clifford Nass. 1999. Effects of humour in task-oriented human-computer interaction and computer-mediated communication: A direct test of SRCT theory. *Human-Computer Interaction* 14(4). 395–435.
- Morreall, John. 1997. *Humor works*. Amherst: Human Resource Development Press.
- Morreall, John. 2020. Philosophy of humor. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Available at: <https://plato.stanford.edu/archives/fall2020/entries/humor/>.
- Niewiadomski, Radosław, Jennifer Hofmann, Jérôme Urbain, Tracey Platt, Johannes Wagner & Bilal Piot. 2013. Laugh-aware virtual agent and its impact on user amusement. In *Proceedings of 12th International Conference on Autonomous Agents and Multi-Agent Systems*, 619–626.
- Nijholt, Anton. 2002. Embodied agents: A new impetus to humor research. In *Proceedings of Twente Workshop on Language Technology*, vol. 20, 101–112. Enschede: University of Twente.
- Nijholt, Anton, Andreea Niculescu, Alessandro Valitutti & Rafael E. Banchs. 2017. Humor in human-computer interaction: A short survey. In *INTERACT 2017, At Mumbai, India*, adjunct proceedings.
- Nijholt, Anton. 2018. Robotic stand-up comedy: State-of-the-art. In N. Streitz & S. Konomi (eds.), *Distributed, Ambient and Pervasive Interactions, 6th International Conference, DAPI 2018*, LNCS 10921, 391–410. Springer.
- Oliveira, Hugo Gonçalves, Diogo Costa & Alexandre Miguel Pinto. 2016 One does not simply produce funny memes! –explorations on the automatic generation of internet humor. *Proceedings of the Seventh International Conference on Computational Creativity (ICCC 2016), Paris, France*, 238–245.
- Pezzullo, Lydia G., Joseph B. Wiggins, Megan H. Frankosky, Wookhee Min, Kristy Elizabeth Boyer, Bradford W. Mott, Eric N. Wiebe & James C. Lester. 2017. “Thanks Alisha, keep in touch”: Gender effects and engagement with virtual learning companions. In *International Conference on Artificial Intelligence in Education*, 299–310. Cham, Switzerland: Springer.
- Piot, Bilal, Pietquin Olivier & Matthieu Geist. 2014. Predicting when to laugh with structured classification. In *Proceedings of Interspeech 2014*, 1786–1790. Singapore.

- Radev, Dragomir, Amanda Stent, Joel Tetreault, Aasish Pappu, Aikaterini Iliakopoulou, Agustin Chanfreau, Paloma de Juan, Jordi Vallmitjana, Alejandro Jaimes, Rahul Jha & Bob Mankoff. 2015. Humor in collective discourse: Unsupervised funniness detection in the New Yorker cartoon caption contest. *arXiv preprint arXiv:1506.08126*.
- Raskin, Victor. 2002. Quo vadis computational humor? In *Proceedings of Twente Workshop on Language Technology*, vol. 20, 31–46. Enschede: University of Twente.
- Ritchie, Graeme. 2003. The JAPE riddle generator: Technical specification. *Informatics Research Report EDI-INF-RR-0158*.
- Ritschel, Hannes, Ilhan Aslan, David Sedlbauer & Elisabeth André. 2019. Irony man: Augmenting a social robot with the ability to use irony in multimodal communication with humans. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 86–94.
- Romano, Daniela. 2007. The look, the emotion, the language and the behaviour of a companion at real-time. In *Artificial Companions in Society: Perspectives on the Present and Future*, vol. 38–40. Oxford: Oxford Internet Institute.
- Ruch, Willibald. 2002. Computers with a personality? Lessons to be learned from studies of the psychology of humor. In *Proceedings of Twente Workshop on Language Technology*, vol. 20, 57–70. Enschede: University of Twente.
- Ruch, Willibald, Salvatore Attardo & Victor Raskin. 1993. Towards an empirical verification of the general theory of verbal humor. *Humor* 6(2). 123–136.
- Rychlowska, Magdalena, Gary McKeown, Ian Sneddon & William Curran. 2019. Not only decibels: Exploring human judgments of laughter intensity. In Jonathan Ginzburg & Catherine Pelachaud (eds.), *Proceedings of Laughter Workshop*, 30–33. Paris: Sorbonne Université.
- Stock, Oliviero & Carlo Strapparava. 2002. HAHAAcronym: Humorous agents for humorous acronyms. In *Proceedings of Twente Workshop on Language Technology*, vol. 20, 125–136. Enschede: University of Twente.
- Tinholt, Hans Wim & Anton Nijholt. 2007. Computational humour: Utilizing cross-reference ambiguity for conversational jokes. In *International Workshop on Fuzzy Logic and Applications*, 477–483. Berlin, Heidelberg: Springer.
- Torre, Ilaria, Emma Carrigan, Killian McCabe, Rachel McDonnell & Naomi Harte. 2019. Cooperating with a smiling avatar: When face and voice matter. In Jonathan Ginzburg & Catherine Pelachaud (eds.), *Proceedings of Laughter Workshop*, vol. 45. Paris: Sorbonne Université.
- Umetani, T., A. Nadamoto & T. Kitamura. 2017. Manzai robots: Entertainment robots as passive media based on autocreated manzai scripts from web news articles. In R. Nakatsu, M. Rauterberg & P. Ciancarini (eds.), *Handbook of digital games and entertainment technologies*, 1041–1068. Singapore: Springer.
- Valitutti, Alessandro. 2009. Computational production of affect-based verbal humorous expressions. PhD dissertation. University of Trento.
- Valitutti, Alessandro & Tony Veale. 2015. Inducing an ironic effect in automated tweets. In *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)*, 153–159. IEEE.
- Weber, Klaus, Hannes Ritschel, Ilhan Aslan, Florian Lingenfelser & Elisabeth André. 2018. How to shape the humor of a robot – social behavior adaptation based on reinforcement learning. In *Proceedings of 20th ACM International Conference on Multimodal Interaction*, 154–162.
- Wong, Josiah & Avelino J. Gonzalez. 2016. CATCH: Injecting “contextual and timely conversational humor” into lifelike avatars. In *The Twenty-Ninth International Flairs Conference*, 134–139.

Bionote

Roddy Cowie

Queen's University, Belfast, Northern Ireland

r.cowie@qub.ac.uk

<https://orcid.org/0000-0003-3480-2223>

Roddy Cowie is an emeritus professor of psychology at Queen's University, Belfast. He was a pioneer in research on affective computing, and is a fellow of the Association for the Advancement of Affective Computing.