P&D Philosophy & Digitality
An Open Access, Peer-Reviewed Journal

On the history of affective computing and the necessity for positive computing

A Discussion with Rafael Calvo and Anna Tuschling/Eva Weber-Guskar as interviewer.

Rafael Calvo, Anna Tuschling, Eva Weber-Guskar

Abstract: Rafael Calvo assesses in a conversation with Anna Tuschling and Eva Weber-Guskar, how the approach has evolved and what new developments he has observed. He elaborates on his concept of positive computing, which is guided by ethical boundaries and carefully evaluates collaborations with economic actors. Calvo advocates for a scientific form of affective computing that is not bound by economic interests and always takes the well-being of users into account. Furthermore, he attributes the interdisciplinary and public scepticism towards affective computing and emotion AI partly to a misunderstanding of fundamental assumptions. More specifically, he considers those criticisms to be misguided which are based on an overestimation of the role of reductionist conceptions of emotion in affective computing.

Rafael Calvo is a pioneer in the field of affective computing, having helped establish and advance the field through his groundbreaking publications and as editor of the acclaimed *Handbook of Affective Computing*. His work brings together insights from engineering, psychological theory, ethics in AI, and computer design, with a particular focus on learning, well-being, and health. Rafael Calvo is Chair in Engineering Design at the Dyson School of Design Engineering at Imperial College London: https://profiles.imperial.ac.uk/r.calvo.

Keywords: positive computing; history of affective computing; basic emotions; Emotion AI; Human–computer interaction

1 **Anna Tuschling/Eva Weber-Guskar:** How did the field of affective computing develop during the last decades? You've built the research field of affective computing together with others. We are curious what your personal view is: What do you consider milestones? And what seemed promising but didn't take off?

2 **Rafael Calvo:** I guess one of the original key publications, sometimes pointed out as the original one, is Rosalyn Picard's book on affective computing (Picard 1997). It just happened that I was on sabbatical that I came upon affective computing. I had done several sabbaticals at the University of Cambridge and they have always been a way to refocus my research. I would take six months or a year to find a new research focus. And at that point in time, I was trying to come back to psychology. It must have been 2007 onwards. At that point, I was spending a lot of time reading about emotions in particular. I was reading Lord Sherrington who was describing a study on emotions in monkeys (Sherrington 1900) and this historical work got me interested. In a later sabbatical, that was 2015, the Oxford Handbook on Affective Computing came out (Calvo et al. 2014). And we also, at the same time, or a couple of years earlier, had started the IEEE Transactions on Affective Computing (Institute of Electrical and Electronics Engineers 2010). So I had the first article on the topic of Affective Computing in the inaugural issue (Calvo/D'Mello 2010).

3 Rosalind Picard's book postulated that we had to work on all the three areas we envisioned for affective computing: Emotion Detection, Emotion Understanding, and Emotion Generation. We contributed to Emotion Detection and to Emotion Understanding. In a way, we worked on Emotion Generation as well or about computers feeling emotions. There's a chapter about that in her book, but that soon dropped out of the original

Contact: Rafael Calvo <r.calvo@imperial.ac.uk>, Anna Tuschling <anna.tuschling@rub.de>, Eva Weber-Guskar <eva.weber-guskar@rub.de>
https://doi.org/10.18716/ojs/pd.v2i2.14039

research. I discussed it with her later, and most of us gave up on emotion generation—this might have changed with the new advances in AI. The mainstream focus appeared to be on the basic emotions work that Ekman and others had done (Ekman 1978). But interestingly, while Emotion Detection remains the biggest area, Ekman's model was one of the very early things that were dropped. For different reasons: One of them is because those are the emotions that don't appear in human-computer interaction (HCI). We don't normally experience these emotions when we are sitting in front of a computer. It's sadness or happiness or others that are not emotions that will come up like that. But interestingly, that's what most of the people that are critical of affective computing have picked up on, that is the methodological concerns about the basic emotions.

4 We worked a lot at the beginning on looking at the models of representation. So when you're working on AI, machine learning, all the problems are really about representation. You have an input and an output, and you have to represent knowledge that comes in and knowledge that comes out, and then you have the task of just mapping one to the other. It's quite a simple thing, but all of machine learning is just about that. It's the machine learning to pattern and to map from here to here. But of course, this has deep implications aside from the purely computational work. When you get deeper and deeper, knowledge representation has ethical implications, has psychological implications, because of tons of different issues and challenges.

5 So in this problem you can say, well, I do facial recognition. The input is my facial expressions. And then the output is its basic emotions. Well maybe it's an interesting mathematical problem. But those emotions are not useful. They—as I said—don't happen very often in computer use. So I think, very soon people started giving up on that. What is interesting is that also you can represent the output as categories. So these are categorical models. In the categorical models you can come up with many other lists of emotions. So very soon after, people started looking for other categories that could be more useful. So we did work on boredom and engagement.

6 **Anna Tuschling/Eva Weber-Guskar:** Would you agree that there are specific computer-related emotions that are originally embedded in human-computer interactions? And do you think they became more interesting for you than basic emotions? You worked a lot on learning **and** on how computers can actually tutor: Are these "HCI" emotions related only to learning environments?

7 **Rafael Calvo**: One of my colleagues, Sidney D'Mello, worked a lot on confusion and the value of detecting and then triggering confusion. And he shows how in e-learning environments confusion could be a very useful technique for supporting learning because it helps people question some of the underlying concepts. So, in any learning environment I can provide you with information that confuses you, and then you start questioning the process. This can be useful in many contexts. Or take repetition and boredom. If a learning system produces boredom, the system can adapt itself to be more challenging, and therefore could become more engaging. Or if it's too challenging, too difficult, the system can provide more feedback.

8 **Anna Tuschling/Eva Weber-Guskar:** The concept of basic emotions developed by Paul Ekman and Wallace Friesen with its focus on less than ten discrete emotions seemed like a perfect match for its adaption in affective computing. Why was it dropped, in your opinion, long before the profound critique by Lisa Feldman Barrett (Barrett 2006; Barrett et al. 2019)? Was it not useful for the pattern generation that you mentioned?

9 **Rafael Calvo:** To answer this question, I have to come back to my earlier thoughts with the categorial model. So we have two models of emotions and the categorical model is one of them. And one of the exercises we do is about changing which categories are useful in a particular use case and in a particular context. The other one is the dimensional model. And we have had 2 or 3 different dimensional models. Well, I think that differentiation is something that Feldman Barrett doesn't go too much into.

10 But she makes one thing very clear and I agree completely with her: what we have (in AC) are these constructs of affect and emotion that are categories and the labels that humans have, and each individual has learned to associate an emotion to certain internal states, and sometimes it can be very simplistic to assume that everybody will assign the same label to the same internal state. There are many layers of complexity in this field. One way we address these challenges is through dimensional models, which help us tackle some of these concerns in various situations. Another point worth noting is the diversity of disciplines and people represented in Affective Computing research. For example, if you look at publications like *IEEE Transactions*, you'll find contributors from a wide range of backgrounds. Some authors are deeply rooted in machine learning, focus-

ing on the mathematical and computer science aspects. Other colleagues come from human-computer interaction, where the emphasis shifts toward the psychological and user-centered dimensions. The discipline of HCI itself is quite broad, and the approaches can differ significantly depending on the perspective. In my own academic journey, I've experienced this firsthand. My master's thesis, for instance, focused on neural networks and their application to predicting sunspot cycles—an 11-year pattern. Back in the early 1990s, much of my work involved pattern matching, comparing neural network approaches to more traditional mathematical methods. From that technical, computer science driven perspective, knowledge representation might not seem particularly meaningful. However, as I've transitioned more toward psychology—much like others in the field, such as Rosalind Picard—I've come to see how the focus shifts. What matters most can change dramatically depending on whether you're approaching the problem from a computer science or a human-computer interaction standpoint.

11 **Anna Tuschling/Eva Weber-Guskar:** We would like to come back to the other distinction you made at the beginning. You have distinguished three different forms of affective computing, first Emotion Detection, second Emotion Understanding—as you call it—and third Emotion Generation. You told us, that you focused your research on the first two forms of affective computing. Can you give us more examples from your various research fields?

12 **Rafael Calvo:** Yeah. Let me see if I find so. One of the relevant, interesting studies was a project to help medical students learn communication skills (Liu et al. 2016). That was, I think, quite a successful project that I worked on when I was at the University of Sydney. I developed a platform designed to teach and enhance communication skills, leveraging teleconferencing technology. The platform facilitates live, interactive interviews between students and simulated interviewees (such as actors or volunteers), with real-time feedback and session recordings available for review. Each clinic within the platform is built with a user-centered design, carefully considering the needs of all stakeholders as students, simulated volunteers, actors, teachers, and academic institutions. To access the platform, users receive a role-specific link. From there, the entire process is streamlined through an automated system, guided by a simple, instruction-based interface. Registration for each clinic is completed online, enabling efficient data collection. Once registered, users are directed to a homepage that provides all necessary resources. The platform also includes built-in training modules, which participants must complete before engaging in any further activities. After preparation, participants can organize appointments directly through an automated calendar system, which displays availability and allows for seamless scheduling. This platform was originally created to address a critical need: training psychology and medical students in affective patient communication. While the concept was developed before the COVID-19 pandemic, its relevance has only grown as online communication has become ubiquitous. For decades, hundreds of students annually used the platform, allowing universities to transition much of their communication training online. This shift reduced the logistical challenges of coordinating actors, students, and physical spaces, making training more accessible and scalable.

13 Initially, we explored the use of affective computing and specifically Emotion Detection to provide students with actionable feedback. Affective and effective communication is vital for doctors, as misunderstandings or misinterpretations of patient emotions can hinder care. While detecting complex emotions like happiness or sadness proved less useful, we focused on simpler, observable cues: speech rate, volume, pitch, smiling, and frowning. These metrics, though not labeled as emotions, offer concrete insights. For example, the platform can detect and highlight moments when a student smiles inappropriately, allowing actors to point out these instances and students to review and reflect on their behavior. We conducted a randomized controlled trial with hundreds of students, demonstrating that these tools significantly improve communication skills. Students often remain unaware of nonverbal cues, such as inappropriate smiling or mismatched facial expressions, which can impact their effectiveness. Mimicry such as nodding or mirroring body language was also emphasized, as it fosters rapport and trust. While this approach may not align with traditional affective computing, it employs similar techniques to recognize and categorize behaviors that are contextually valuable. Many in the field have since adopted comparable methods, moving beyond emotion labeling to focus on actionable, context-specific feedback.

14 **Anna Tuschling/Eva Weber-Guskar:** How did your view of human emotions change due to your work on affective computing and in what way? You told us, you became more and more interested in psychological questions and research? Did this also change your view of what emotions are?

15 **Rafael Calvo:** Yeah, I think that's a very important fact because for psychologists like Feldman, maybe the discipline is just about recognizing a set of categories that a psychologist will call emotions and their emotional processes. In affective computing, it's basically just recognizing mental states.

16 Over the decades, I think more and more people have moved away from a strict concept of emotions and started thinking more broadly in terms of mental states. Going back to your question about whether my academic work has changed my views on emotions, I'd actually say that the more interesting answer is that it wasn't my machine learning research—it was my work on AI ethics that changed my perspective. In my machine learning work, I've always been careful not to describe these systems as having emotions. That's why we deliberately use terms like *smiling*, *nodding*, or *head shaking* rather than claiming that a system is *happy* or *sad*. Nobody assumes that when a doctor smiles it necessarily means, the doctor is happy. Similarly, we're not suggesting that a machine displaying a smile is experiencing happiness. Some critics of affective computing have implied that researchers in the field believe machines genuinely have emotions, but I think that's a misrepresentation. In our work, we intentionally avoid making direct connections between observable behaviors and emotional experiences. I think that's fairly common practice. What has really changed my perspective is a growing awareness of AI ethics. It made me realize that knowledge representation is itself a deeply ethical issue, and one that, in my view, neither psychologists nor AI ethicists have examined deeply enough.

17 I think Lisa Feldman Barrett's work is excellent. I really enjoyed her book (Barrett 2017), and I agree with many of her arguments. However, one point where I would differ is that whether an error margin is "large" or "small" is entirely dependent on the application domain.

18 Take emotion recognition, for example. If you achieve only 5 % better than random accuracy in a psychological setting, that's a very poor result. A clinical psychologist who could recognize people's emotions only slightly better than chance would not be considered competent. But for a computer system, the context is very different. That's where my perspective shifted. The problem is that affective computing is increasingly being deployed in domains like marketing and behavioral influence—applications that are quite different from the ones many of us originally envisioned. In those contexts, a 5 % improvement can be enormously valuable. If an algorithm increases Amazon's or Google's revenues by 5 %, you're talking about hundreds of millions of dollars. Even a seemingly modest increase in the accuracy of a marketing algorithm can have a massive commercial impact.

19 I think, that's something we still don't discuss enough. I haven't heard many colleagues in psychology or philosophy address this point, but I believe it's crucial. If we want to make meaningful critiques of these technologies, we need to understand the business contexts in which they're actually being deployed. In the same way that I think engineers need to learn more about ethics, and who really benefits from affective computing (i.e. large marketing companies), I also believe that ethicists and psychologists need to learn more about technology—not just how it works, but also about the business contexts in which it is developed and deployed. Without that broader understanding, it's difficult to appreciate where these technologies are actually being used and what incentives are driving their adoption. Otherwise, our attempts to inform government policy or shape public debate can come across as rather naive.

20 **Anna Tuschling/Eva Weber-Guskar:** A lot of behavioral research is conducted at major tech companies. That's exactly what you're criticizing as a problem in your statement. Do you have any experience with any of the research labs at major tech companies?

21 **Rafael Calvo:** No, not generally. We have collaborations—or at least contacts. We've tried multiple times to establish collaborations with people at Microsoft. Mary Czerwinski, for example, is a very well-known affective computing researcher at Microsoft. But generally, once things move out of the research domain, they become part of the company's industrial or commercial interests, and it becomes harder to collaborate on those kinds of projects. Personally, as I started seeing the negative side of these developments, I became less interested in collaborating on them. I suppose that's my perspective as someone interested in AI ethics.

22 **Anna Tuschling/Eva Weber-Guskar:** What is your opinion on the EU AI Act's restrictions on Emotion Detection and biometric systems, particularly regarding their impact on research and practical applications?

23 **Rafael Calvo:** When there is a high-risk situation, I think it is important to legislate and manage that risk. We have not adequately acknowledged the risks associated with these kinds of tools, particularly when it comes to manipulation. I believe, manipulation has a much deeper impact on who we are than we often realize. I'm not just referring to the economic or political aspects, such as the surveillance economy, but to the long-term effects on individuals and society. As these technologies become more effective, they can reduce our sense of autonomy, which ultimately affects our well-being. My research has evolved from machine learning to affective computing, at the intersection of machine learning, human-computer interaction, and psychology. More recently, I have focused on AI for mental health and supporting well-being.

24 The framework I use is **Self-Determination Theory**, which identifies three basic psychological needs that are essential for human flourishing: autonomy, competence, and relatedness. I think, manipulation technologies, including affective computing, can undermine people's sense of autonomy. In that sense, I agree with the EU AI Act (2024) in classifying these technologies as high risk, particularly in areas such as education.

25 I do see this as a significant concern. When deploying AI interventions with children, we should apply the same level of caution that we would when introducing interventions for people with health conditions. Even if children are healthy, they are still a vulnerable population and deserve careful protection.

26 **Anna Tuschling/Eva Weber-Guskar:** So you would say in certain instances, affective computing can become manipulative?

27 **Refael Calvo:** Yes, definitely. The technology industry is largely driven by advertising. For example, a major source of Google's and other's revenue comes from marketing. If companies could increase the accuracy of their advertisements by even 3 % using affective computing techniques, many would see that as a significant business opportunity.

28 However, I see it differently. To me, this represents more effective manipulation, which I believe is harmful in the long term. Education is a different context. We have to be careful with the wording, but education is intentionally about helping people change. Its purpose is to support individuals in developing and growing in ways that benefit them, while giving them the tools to direct that change themselves. What I would not trust is a technology company whose profit depends on influencing that change. That creates a conflict of interest and raises serious ethical concerns.

29 **Anna Tuschling/Eva Weber-Guskar:** You raised several important—and in some ways troubling—points. If we shift the perspective to the computer, could one also argue that affective computing helps computers understand how humans feel and interact with one another? In the models you mentioned, we provide computers with information about human behavior and emotions. In your opinion, what role does "affective computing" play in enabling computers to interpret and model these human interactions?

30 **Rafael Calvo:** Yeah, but let me clarify, because I think that's a common misconception. It took me a while to fully understand this myself: the computer is not a value-neutral object. It is not simply an impartial tool. Behind every computer, and the systems it runs, there are companies, institutions, and groups with particular interests and priorities. It's not just a philosophical or abstract object—it is also a commercial one. Once you become aware of that, your perception of technology changes.

31 **Anna Tuschling/Eva Weber-Guskar:** As affective computing is increasingly combined with robotics and large language models to create more comprehensive and human-like AI systems, what normative and ethical issues do you see arising from this convergence?

32 **Rafael Calvo:** Well, large language models, like affective computing at one level, are essentially pattern-matching systems. You provide a sequence of input tokens, and the model predicts the next token. The key limitation is that these models are not grounded in embodied cognition—they do not have a body or direct experience of the world. Many of the problems associated with large language models stem from this lack of embodied knowledge. Fundamentally, they are predicting one token after another based on statistical patterns. In some respects, affective computing and Emotion Recognition face a similar challenge. They, too, rely on pattern

matching rather than embodied experience. With enough data, these systems can appear highly accurate in a large percentage of cases. However, they can also produce unexpected or inappropriate outputs that seem completely out of context. That happens because they lack the embodied understanding that underlies human emotions.

33 Emotions are deeply connected to embodied cognition, and that is a major element missing from these systems. When we are talking about language models alone, this limitation is particularly evident. There are newer approaches that apply transformer models to robotics, and some researchers argue that this represents a fundamentally different direction because the models become connected to an embodied agent interacting with the physical world. But when we are talking solely about language models, they remain, at their core, systems that predict sequences of tokens. The absence of embodiment makes a significant difference.

34 **Anna Tuschling/Eva Weber-Guskar:** Some colleagues claim that technology could become sentient (Chalmers 2023). Can a holistic system combining advanced sensors, affective computing, and language models develop its own role models or world representation, and could this integration lead to truly new, emergent capabilities?

35 **Rafael Calvo:** No, I don't think these systems will become sentient, at least not in the sense in which we usually understand sentience. In my view, that question is not really appropriate until machines have embodied cognition, embodied understanding, and embodied intelligence. For me, sentience requires a body. I know many of my colleagues would disagree, but until that changes, what we will have are increasingly sophisticated systems that are better at pattern recognition. I can imagine how affective computing could improve some robot's performance. If the robot detects that I am smiling, it could use that information to generate a better prompt for a large language model such as ChatGPT. Rather than simply asking for "a nice story," it could tailor the prompt to my emotional state. If I appear sad, it might ask for a story intended to lift my mood. If I appear happy, it could generate a story that reinforces or celebrates that feeling. In that sense, affective computing can make the interaction appear much more natural. However, this does not mean the robot actually feels anything. It simply becomes a more convincing pretender.

36 **Anna Tuschling/Eva Weber-Guskar:** You place a great deal of emphasis on embodiment. Earlier, you also mentioned that your research has become increasingly focused on human-computer interaction. At the moment, however, most interactions are still screen-based. As in our conversation today, the interaction is largely limited to the upper body—primarily the face, which is one of the most expressive parts of the human body.

37 One could therefore imagine a different kind of interface. Robots equipped with multiple sensors might enable a much richer form of human-computer interaction by taking into account not only facial expressions but also posture, gait, gestures, and the way people move through space. These are all important components of emotional expression.

38 **Rafael Calvo:** Exactly. My move toward human-computer interaction and psychology was motivated by a desire to better understand human well-being. As I argued in *Positive Computing* (Calvo/Peters 2014), understanding emotions alone is neither necessary nor sufficient for supporting well-being. My broader view is that all technologies, not just computers, should be designed to support human well-being. I don't see much value in developing technologies that fail to do so. Human well-being depends on a number of basic psychological needs, alongside physical and other needs, and technology should be designed to support these rather than undermine them. If a technology hinders these fundamental needs, then perhaps it should not exist in the first place. This is where I differ from some colleagues in human-computer interaction and technology, who sometimes view technology as the solution to every problem. In many situations, I don't think that is the right approach.

39 Of course, it would be wonderful if technology allowed the three of us to sit together naturally in the same virtual space. Some of the latest developments are moving in that direction. One of my students, for example, is working on Apple's Vision Pro and spatial computing. Even so, we have to be extremely careful, because different technologies are driven by very different business models. If you compare companies such as Meta with Apple's approach to spatial computing, there are important philosophical differences. Meta's vision is centered on avatars—creating an alternative digital representation of yourself. Apple's approach, at least conceptually, is

more about placing you in another location while preserving your own identity rather than replacing it with an avatar.

40 **Anna Tuschling/Eva Weber-Guskar:** That's an interesting distinction. Apple deliberately uses the term *spatial computing* rather than *virtual reality*, which reflects a different way of thinking about the technology. Even so, many questions remain—for example, whether these devices are appropriate for use in public spaces, and whether increasing bodily representation in digital communication is actually desirable.

41 **Rafael Calvo:** Exactly. We have to anticipate unintended consequences. That's something I increasingly emphasize in my teaching. We need to help engineers and designers think systematically about the possible consequences of the technologies they create. Meta's vision for the glasses is all about mass surveillance and us collecting data for them to use in their algorithms.

42 Apple's approach may be preferable to one focused entirely on surveillance, or avatars and idealized digital identities, but it still carries significant risks. Once you create a detailed digital model of a person's body, that representation can be copied, manipulated, or misused. There are obvious commercial incentives to exploit such models, and we need to think carefully about those possibilities. Personally, I would be deeply concerned if someone could take a realistic digital model of me and use it for purposes I never consented to. The potential unintended consequences are substantial and, frankly, quite alarming.

43 This is why engineers need to become much better at anticipating these risks. At the same time, philosophers and ethicists often remain too far removed from the technical details to fully understand what is possible. We therefore need genuine collaboration between technical and humanistic disciplines if we are to anticipate and address these challenges effectively.

References

Barrett, Lisa Feldman. 2006. "Are Emotions Natural Kinds?" *Perspectives on Psychological Science* 1 (1): 28–58. https://doi.org/10.1111/j.1745-6916.2006.00003.x.

Barrett, Lisa Feldman. 2017. *How Emotions Are Made: The Secret Life of the Brain*. Houghton Mifflin Harcourt.

Barrett, Lisa Feldman, Ralph Adolphs, Stacy Marsella, Aleix M. Martinez, and Seth D. Pollak. 2019. "Emotional Expressions Reconsidered: Challenges to Inferring Emotion From Human Facial Movements." *Psychological Science in the Public Interest* 20 (1): 1–68. https://doi.org/10.1177/1529100619832930.

Calvo, Rafael A., and Dorian Peters. 2014. *Positive Computing: Technology for Wellbeing and Human Potential*. MIT Press.

Calvo, Rafael, and Sidney D'Mello. 2010. "Affect Detection: An Interdisciplinary Review of Models, Methods, and Their Applications." *IEEE Transactions on Affective Computing* 1 (1): 18–37.

Calvo, Rafael, Sidney D'Mello, Jonathan Gratch, and Arvid Kappas, eds. 2014. *The Oxford Handbook of Affective Computing*. Oxford University Press.

Chalmers, David. 2023. "Could a Large Language Model Be Conscious?" *Boston Review*. https://www.bostonreview.net/articles/could-a-large-language-model-be-conscious/.

Ekman, Paul. 1978. *Facial Action Coding System*. Consulting Psychologists Press.

Institute of Electrical and Electronics Engineers. 2010. *IEEE Transactions on Affective Computing*. IEEE.

Liu, Chunfeng, Renee L. Lim, Kathryn L. McCabe, Silas Taylor, and Rafael A. Calvo. 2016. "A Web-Based Telehealth Training Platform Incorporating Automated Nonverbal Behavior Feedback for Teaching Communication Skills to Medical Students: A Randomized Crossover Study." *Journal of Medical Internet Research* 18 (9): e246. https://doi.org/10.2196/jmir.6299.

Picard, Rosalind. 1997. *Affective Computing*. MIT Press.

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA Relevance) (2024). http://data.europa.eu/eli/reg/2024/1689/oj.

Sherrington, C. S. 1900. "Experimentation on Emotion." *Nature* 62 (1605): 328–31. https://doi.org/10.1038/062328a0.