

The Joi of holograms

PAUL SMART

Introduction

AS A FORM OF HOLOGRAMMATIC artificial intelligence (AI), Joi is a new addition to the *Blade Runner* universe. We know, of course, that Joi is synthetic, in the sense of being a technological artefact, but as the movie progresses it becomes increasingly difficult to see her as anything other than a virtual person. By the end of the movie, we are left with a number of seemingly unanswered questions: Who or what was Joi? Was she a sentient individual, capable of subjective experience? Was her love for K genuine? Did she qualify as a person, on a par, perhaps, with replicants and humans? Or was she nothing more than a computational ruse, an ethereal play of light and sound intended to obscure the grim reality of K's rather joyless existence?

Such questions identify some of the points of philosophical interest raised by Joi. But Joi also raises issues of a more metaphysical nature. Her hologrammatic status, for example, encourages us to reflect on the philosophical significance of a number of new technologies (e.g. Microsoft's HoloLens) and consider their implications for the scope of what is sometimes called "cinematic philosophy" (i.e., the philosophical study of the cinematic medium). As we will see, Joi presents us with some familiar philosophical puzzles and problems, but she also serves

as the means by which a number of new ideas and issues are brought to light.

The enigmatic hologram

One of the issues raised by *Joi* concerns her status as a sentient being. Let us call this *the sentience issue*:

The Sentience Issue

Is *Joi* a sentient individual, with an inner mental life similar to our own? Are her actions suggestive of a capacity to experience pain, hope, joy, and loss? Or do they reflect nothing more than the instantiation of a particularly well-crafted and (it has to be said) endearing computational algorithm?

From a philosophical perspective, the sentience issue is perhaps best understood as a claim about phenomenal consciousness. To consider the sentience issue is thus to entertain the possibility that there is, in the words of Thomas Nagel (1974), “something it is like” to be *Joi*. *Joi* certainly behaves in a way that is consistent with a positive response to the sentience issue. But does this mean she qualifies as a sentient individual, with all the thoughts, feelings, and emotions that we ascribe to a typical human being? Is she a being who is capable of feeling all the things that we (as humans) might experience while watching *Blade Runner 2049*? Or is she no more sentient than the play of light that we see projected on the “big screen?”

Before going further, we ought to ask ourselves why issues of sentient standing arise in the case of *Joi*. Why, in particular, are we inclined to question *Joi*’s status as a sentient individual when we don’t (for the most part) ask similar questions about *K*, *Deckard*, or anyone else in the movie?

One answer to this question relates to *Joi*’s hologrammatic status—the fact that she exists as an insubstantial, ghost-like entity. To be sure, *Joi* is a rather unusual kind of AI entity; nevertheless, I very much doubt that *Joi*’s status as a hologram can tell us very much about her experiential capabilities. Perhaps it is true, for example, that *Joi* is, in some sense, more “soul than substance.” But even if we are prepared to accept this Cartesian distinction between mind (soul) and matter (substance), why

should Joi's seeming ethereality undermine her capacity for phenomenal consciousness? After all, it wasn't so long ago that philosophers deemed mechanical robots to be devoid of experience on account of the fact that they lacked a soul!

A second (and, in my view, more important) reason to question Joi's status as a sentient individual relates to the way she is presented in the context of the movie. Consider, for example, the "Joi in rain" scene, where we see Joi venture outside of K's apartment for the first time. At the beginning of this scene, Joi responds to the falling rain in a way that is suggestive of the experience of pleasure. Later, however, as Joi and K exchange a "virtual" kiss, an incoming call (from Lt. Joshi) causes Joi's image to be frozen. The effect is similar to the freeze-framing of a conventional motion picture or movie, and it serves as a striking contrast to the dynamic, animated "image" we had of Joi only a few seconds earlier. In particular, we are reminded of the potential parallels between Joi and the cinematic medium itself: If Joi is the sort of thing that can exist as a freeze-frame shot, then perhaps she is no more "real" or sentient than is the very movie in which she appears.

The Joi in rain scene thus leads to a tension in the way we think about Joi. On the one hand, the richness and expressivity of Joi's behaviour encourages us to view her as a sentient individual. On the other hand, however, we are presented with a state-of-affairs (i.e., the freeze-frame) that makes us doubt the evidence of our senses. If, for example, we see consciousness as a process, then it seems unlikely that Joi could be conscious when she is in freeze-frame mode. But, at the same time, Joi may be like many other forms of moving image (or motion picture), in the sense that she is rendered as a series of still images that are presented in quick succession, so as to give the "illusion" of continuous movement. If this is so—if Joi is, in fact, nothing more than a succession of still images, each of which is effectively frozen in time—then why should we see the composite of those static snapshots as giving rise to something as genuinely sentient, as opposed to something that merely looks like the real deal?

There are, therefore, a number of reasons why issues of sentience are particularly pertinent in the case of Joi. Yet, despite the fact that the movie encourages us to confront such issues (via the use of cinematic ploys like the aforementioned freeze-frame shot), there is, as far as I can tell, nothing in *Blade Runner 2049* that would help us to resolve the

sentience issue. Perhaps, for example, it is the case that Joi qualifies as a cinematic resource. But why should that resolve the sentience issue? Doesn't it merely beg the question as to why we think the cinematic medium is unable to serve as a medium in which synthetic forms of conscious experience might (one day) be realised?

With regard to the sentience issue, then, Joi is an enigma, plain and simple. And that may very well be the end of the matter. For no amount of probing or analysis of the movie's content is, I think, likely to yield some new insight that settles the sentience issue one way or the other. There is, I suggest, no (hidden) fact to be discovered in the movie, and to subject the movie to further scrutiny in the hope of finding such a fact is probably a futile exercise. For better or worse, then, the sentience issue is likely to remain unresolved: Joi's status as a sentient being is not something that we can discover from watching the movie (or subjecting it to detailed analysis).

Should we worry about this rather unfortunate state-of-affairs? Does Joi's enigmatic status in respect of the sentience issue reduce her philosophical significance, or in any way undermine the philosophical credentials of the movie? One reason to think that the answer to these questions is a resounding "no" stems from a consideration of the philosophical literature itself. Of particular interest is the fact that our philosophical (and scientific) understanding of what it means to be a sentient individual—or, indeed, how we might go about resolving issues of sentient standing—is surprisingly poor. There is, in short, no consensus on how issues of sentience ought to be resolved. That does not mean that philosophers have had nothing to say about sentience-related matters (they have!) (e.g. Hyslop, 2018); it simply means that there is, as yet, no definitive answer to the sort of question raised by the sentience issue. From a philosophical standpoint, then, this looks to be one area where cinematic fiction echoes philosophical reality. Indeed, when it comes to issues of sentience, the fan of *Blade Runner 2049* is arguably placed in pretty much the same epistemic position as is the professional philosopher. The professional philosopher is, of course, concerned with a generalised version of the sentience issue, but neither the philosopher nor the moviegoer is presented with any easy answers to what amounts to a common problem. In this sense, *Blade Runner 2049* provides the moviegoer with some insight into the nature of a recognisably important, and yet seemingly intractable, philosophical issue: The movie steadfastly refuses to deliver a definitive answer to

the sentience issue, just as our current philosophical understanding of the world leaves us none the wiser as to how issues of sentient standing ought to be resolved. Perhaps this tells us something about the philosophical value of Joi's enigmatic status in respect of sentience-related issues. If, for example, the philosophical literature tells us that there is no way for issues of sentient standing to be resolved based on an analysis of an individual's behaviour, then why should we expect *Blade Runner 2049* to provide us with a definitive answer to the question of whether or not Joi qualifies as a sentient individual?

Autonomy, authenticity, and virtual love

There is of course much more to Joi than just the sentience issue. Another issue relates to whether or not Joi has any sort of autonomy. Let us call this the *autonomy issue*:

The Autonomy Issue

Does Joi exhibit any sort of autonomy with regard to her behavioural and emotional responses? Is she able to respond in a way that violates the constraints of her programming, or is she forever condemned to obey the computational commands of her creators?

We know that Joi is manufactured by Wallace Corp., and it thus seems reasonable to assume that she comes equipped with some pre-programmed functionality. At the same time, however, it is clear that Joi is a highly adaptive system. As is evidenced by the various Joi-related adverts in the movie, K's Joi is a specific instance of a more general AI product. This means that unless all of Joi's clients are exactly the same as K, which seems unlikely, then she must be able to tailor her responses to suit the demands of specific situations. The question, of course, is whether this sort of adaptive capacity is sufficient for her to evolve beyond the constraints of her programming. Is she able to become a free-thinking, autonomous individual (like K, perhaps), or is she forever condemned to comply with the functional imperatives laid down by those who designed her?

All of this has a bearing on whether or not we regard Joi's love for K as authentic. Joi certainly appears to express affection for K, but does this mean that her love is "real?" Are her responses suggestive of a genuine form of love, or is it merely an ersatz form of love—a "holo"

simulacrum of the real thing? This is yet a further issue that is raised by an analysis of Joi. Let us call it the authenticity issue:

The Authenticity Issue

Is Joi's love for K authentic? Does it (should it) count as "real" or genuine love?

For the purposes of illustration, let us assume that Joi is capable of feeling love, and that she thus counts as a sentient individual (i.e., we adopt a positive response to the aforementioned sentience issue). Does this mean that Joi's love for K is genuine?

One problem we have here is that Joi may not feel love, despite the fact that she is capable of feeling love. She may, in short, be particularly good at "playing the part," that is, a good actor.

But even if Joi *does* feel love, our problems are not quite over; for Joi may have been explicitly programmed to feel love. Remember that we only get to see one instantiation of Joi in the movie. Presumably, however, there are lots of Joi instances in the *Blade Runner* universe, and not all of these will be fortunate enough to end up with a die-hard romantic like K. If Joi has been specifically designed to feel love whenever her owner wants her to feel love, then she will feel love whenever she is required to do so, regardless of who she is with. Clearly, if Joi is compelled to feel love for anyone who wants it, then there is nothing special about Joi's love for K, nor is there anything special about K that makes him particularly deserving of Joi's affections.

We thus confront a puzzle. Although there is a hint in the movie that Joi's "feelings" for K may deepen as the movie progresses,¹ we are never given a straightforward answer to either the autonomy or the authenticity issues. In respect of the authenticity issue, for example, we are never explicitly told that Joi is capable of feeling love; nor are we explicitly told that Joi does indeed love K, even if she is capable of feeling love. It is only when Joi's emanator is about to be crushed by Luv that we hear Joi attempt to proclaim her love for K by saying "I love you." But even here the expression of love is incomplete; the emanator is destroyed before Joi has had a chance to complete her sentence, and the result is not so much "I love you" as it is "I love y—." The result is that we face yet another point of ambiguity that complicates the interpretation of Joi's character.

Let us assume that Joi's love can only count as genuine if she experiences the various emotional highs and lows that we typically

associate with romantic love. That is to say, Joi must *experience* love in order for her love to count as genuine. This implies that a positive response to the sentience issue is a prerequisite for a positive response to the authenticity issue. If Joi is not sentient, then she will be in no better position than if she were sentient but did not feel love, that is, the good actor option in Figure 7.1. Relative to Figure 7.1, then, we have two options, both of which turn on the nature of our response to the autonomy issue: Joi is either a virtual slave (if she has no autonomy) or a virtual lover (if she does). But which is it? Does Joi love K because of who he is, or because of who she is? Does Joi love K merely because she was designed to love K, or does she love him because she chooses to love K? Is it, as Niander Wallace wonders about the nature of Deckard and Rachael's relationship, "Love ... or mathematical precision? Yes? No?"

Wallace's question suggests that we are limited to a binary yes/no response, but I am inclined to think that the actual answer is somewhat less clear-cut. According to one of Joi's advertising slogans, "Joi is

	Feels Love	Does Not Feel Love
Programmed to Love	Virtual Slave	Defective Product
Not Programmed to Love	Virtual Lover	Good Actor

Figure 7.1 Different interpretations of Joi according to whether or not she is programmed to feel love and whether or not she actually feels love. All interpretations assume that Joi is sentient, is capable of experiencing love, and behaves as she does in the movie. Being "Programmed To Love" means that Joi has no capacity for emotional autonomy; she is simply compelled to experience love whenever she is required to do so. "Not Programmed To Love" implies a greater degree of emotional autonomy. Joi might start out as a good actor, for example, but progress to a virtual lover based on the nature of her relationship with a human/replicant client.

whatever you want her to be.” Presumably, this is the sort of functional constraint that lies at the core of Joi’s programming, and it is a constraint that Joi must always obey. It is, let us say, her *prime directive*. Relative to this prime directive, Joi is not autonomous; she is compelled to say and do whatever her owner wants her to say and do. In this sense, then, Joi is a virtual slave (see Figure 7.1).

But what if her owner wasn’t happy with this state-of-affairs? What if Joi’s owner wanted her to be something more than this? “I am so happy to be with you,” says Joi. “You don’t have to say that,” replies K. If Joi learns that K does not want her to be a virtual slave, then is she still compelled to be one? Wouldn’t her failure to transition to a more autonomous state violate the prime directive? And if the prime directive were to be violated, then wouldn’t she count as being autonomous anyway? After all, to remain a virtual slave when K wants her to become a virtual lover is to violate one of the rules (and perhaps the only rule) that prevents her from qualifying as an autonomous entity. In this respect, regardless of whether Joi remains a virtual slave or switches to become a virtual lover, she will count as autonomous if K wants her to be autonomous.

Is there any evidence to suggest that Joi does transition to a more autonomous state? The litmus test, in this case, is whether or not Joi does something that K does not want her to do. According to the prime directive, recall, Joi is required to be whatever we want her to be, but we are also told that Joi is “Everything you want to hear.” Although this sort of constraint appears to be consistent with the prime directive, we might expect that a virtual lover (but not a virtual slave) would have the capacity to say something that K does not want to hear. And there is, indeed, one point in the movie where Joi appears to ask K to do something that he does not want to do. Before they head off to Las Vegas, Joi asks K to delete her from the console and break her antenna:

JOI: You have to delete me from the console.

JOI: My present. Put me there. [Referring to the emanator.]

K: I can’t do that. Think about it: If anything happens to this, that’s it. You’re gone.

JOI: Yes. Like a real girl.

JOI: Please. Joe, please. I want this. But I can’t do it myself.

JOI: Break the antenna.

Inasmuch as we interpret this dialogue as evidence that Joi is saying something that K does not want to hear, then the autonomy issue is, I suggest, resolved in favour of a positive response. If Joi is sentient, then, she begins her time with K as a virtual slave, but she subsequently transitions to a virtual lover based on the nature of her relationship with K (see Figure 7.1).

And as for Wallace's question (Is it love or mathematical precision?), the answer, at least in this case, is probably "both."

Art, ambiguity, and intrinsic value

The way we think about Joi alters our view of events in the movie. Consider, for example, Joi's destruction at the hands (or rather the boot) of Luv. If Joi is neither sentient nor autonomous—if she is something akin to a character in a computer game—then her love for K doesn't appear particularly genuine, her death marks the end of K's digital fantasy, and Luv is guilty of property damage. If, on other hand, Joi is sentient and autonomous—if she is something akin to a virtual person—then her love for K is genuine, her death is a tragedy, and Luv is guilty of murder.

Given that *Blade Runner 2049* is, at root, a work of art, it is perhaps not surprising that it should exhibit some degree of ambiguity. Ambiguity is, of course, a common feature of artistic works, and it has long been seen as relevant to issues of aesthetic evaluation (e.g. Tormey & Tormey, 1983). When it comes to philosophical matters, however, we encounter a problem. The problem is that ambiguity is seldom a celebrated feature of philosophical works, and this leads to a worry about the extent to which movies (*qua* works of art) can be seen to make a philosophical contribution. In particular, if the notion of a philosophical contribution is tied to issues of clarity and precision, and movies (or, at any rate, the philosophically relevant parts of movies) are deemed to be ambiguous, then it is, at best, unclear how a movie could be deemed to make a philosophical contribution (or perhaps qualify as a work of philosophy). In the philosophical literature, this particular worry is nicely captured by the so-called *explicitness objection*. According to Wartenberg (2007, 2009), the explicitness objection highlights the contrast between "the precise and explicit formulation of claims in philosophy texts with the

supposed imprecision and ambiguity of narrative films” (Wartenberg, 2007: 138).

There are a number of ways of responding to the explicitness objection. Wartenberg’s own response is to reject the idea that movies are inherently ambiguous. He distinguishes between the ambiguity of whatever philosophical interpretations are made about a film and the inherent ambiguity of the film itself. According to Wartenberg (2009), the explicitness objection ought to be interpreted as a form of methodological advice: “Provide enough specificity and determinateness in a philosophical interpretation of a film so that the philosophy presented by the film is not inherently ambiguous” (Wartenberg, 2009: 553).

Unfortunately, this particular way of responding to the explicitness objection is unlikely to be of much use when it comes to *Joi*. The reason for this is simple: no amount of precision or exactitude in the philosophical interpretation of *Joi*’s character is likely to alter the fact that she is inherently ambiguous. We can, of course, choose to enforce a particular interpretation for the sake of philosophical analysis (e.g. “*Joi* is sentient”), and there is no reason to think that the philosophical value of whatever claims and arguments are made in respect of that analysis are in any way impugned by the deliberate imposition of one or more interpretational certitudes. The fact remains, however, that we simply aren’t given enough information to determine whether or not *Joi* is (e.g.) a sentient individual. This is simply one area (and there are no doubt others) where the movie fails to yield a determinate response. This has nothing to do with the quality (or otherwise) of the hermeneutic analysis of the movie; it is simply a reflection of the movie’s *inherent* ambiguity when it comes to particular points of philosophical interest.

How, then, ought we respond to the explicitness objection? The response I want to canvass here is what I will dub the *intrinsic value* response. According to this response, the philosophical value of a resource, or its capacity to make a philosophical contribution, is on occasion inextricably tied to the ambiguous nature of the resource itself. To help us understand this claim, consider one of the claims made in respect of the sentience issue: the idea that *Joi*’s enigmatic status parallels the complex, confusing, and uncertain nature of a broader philosophical debate. Just as we may not be able to resolve the sentience issue in the context of the movie, so our capacity to resolve sentience-related issues in the wider world may be similarly constrained. It is in precisely this sense

that ambiguity might be seen to have intrinsic value. For the aim of *Blade Runner 2049* is not to solve the philosophical problem of sentience, or even to illuminate the path to a potential solution. Rather, Joi's enigmatic properties encourage us to confront a philosophical problem and experience some of the doubt, indecision, and ambivalence that accompanies that particular problem. In doing so, the movie helps us understand something about the nature of the problem, such as why it exists, why it is important, why it is difficult to solve and why, ultimately perhaps, it may have no satisfactory solution. There is no sense here in which the movie's inherent ambiguity could be seen to impugn its capacity to make this sort of philosophical(?) contribution. If anything, ambiguity serves as the vehicle through which at least some of the movie's philosophical contributions (and thus its philosophical value) are realised.

Cinematic philosophy beyond the big screen

When we first encounter Joi, we learn that her holographic projection is limited to K's apartment. "How was your day?" asks K. "I'm getting cabin fever," replies Joi.

We later learn that K has purchased a gift for Joi: an emanator. This makes Joi portable. Whereas her projection had been limited to the confines of K's apartment, she is now free to accompany K on his travels and "experience" whatever it is that the world of *Blade Runner 2049* has to offer.

Relative to the movie, of course, the emanator serves as a plot device that enables Joi and K to be "joined at the hip," so to speak, throughout the remainder of the movie (at least until the point that Luv "steps in"). For me, however, the introduction of the emanator marks a point of allegorical significance: it speaks to the way in which emerging technologies, including holographic computing devices, are apt to make virtual cinematic elements (e.g. holograms) a standard feature of physical and social reality. This raises questions about the nature of cinematic philosophy. Should the scope of cinematic philosophy be defined with respect to whatever it is that appears on the big screen (e.g. conventional movies), or should we embrace a more expansive vision of the scope of cinematic philosophy—for example, one that includes the philosophical study of hologrammatic resources? In the same way that Joi's emanator enables her to escape the confines of K's apartment and venture out

into the world, perhaps the emergence of mixed-reality devices calls for cinematic philosophers to take a closer look at the world that lies beyond the borders of the big screen—that is, to embrace the idea that non-film resources should form part of the intellectual remit of cinematic philosophy.

The term *cinematic philosophy* names a specific area of philosophy with an intellectual remit that concerns the philosophical significance and impact of the cinematic medium (e.g. Shamir, 2016; Wartenberg, 2007). For the most part, the scope of this philosophical effort has limited its attention to the sort of resources that we typically associate with the term *cinema*. They include what are termed films, movies, or motion pictures—the sort of things that we typically get to see on the big screen. *Blade Runner 2049* is a perfect example of this sort of cinematic resource, and it is thus a resource that falls within the scope of contemporary cinematic philosophy.

But should the scope of cinematic philosophy be limited to the study of films, movies, motion pictures, and the like? From an etymological perspective, the word “cinema” is derived from the ancient Greek word *kinēma*, meaning motion and movement. This highlights one of the central features of the cinematic medium, namely, its preoccupation with the moving image, the “motion picture.” In a general sense, then, the cinematic medium is the medium of the moving image, and cinematic resources are resources that exploit this medium. Cinematic resources are, in short, images in motion—typically, a series of still images that are presented in rapid succession so as to give the impression of movement.

A Hollywood movie such as *Blade Runner 2049* undoubtedly qualifies as a cinematic resource, but is there any reason to suppose that the extension of the term “cinematic resource” is limited to this type of resource—the sort of resource that appears on the big screen?

It should be relatively clear that the answer to this question is “no.” As noted by Berys Gaut (2010: 1), “Moving images come in many kinds, a fact of which it has been easy to lose sight until recently, given the dominance of traditional photochemical, celluloid-based film for most of the period in which moving images have been subject to theorising.” Gaut (2010) goes on to note the cinematic status of resources such as video games and virtual-reality environments. While these resources have features that are unlike those of a conventional film (e.g. a video

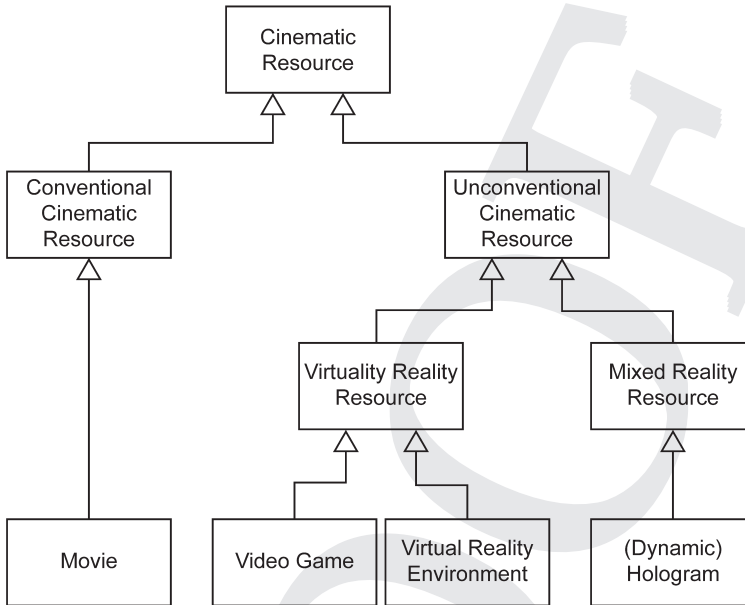


Figure 7.2 Taxonomy of cinematic resources. Triangles symbolise taxonomic or subtype-of relationships. (Note that only dynamic holograms qualify as cinematic resources.)

game is interactive in a way that a conventional film is not), they are nevertheless resources that also rely on the cinematic medium. In this sense, a video game counts as a cinematic resource, just as much as does a Hollywood blockbuster.

Figure 7.2 illustrates a rough (and not necessarily complete) taxonomy of cinematic resources. As can be seen from the figure, cinematic resources are divided into two types: conventional and unconventional cinematic resources. This distinction is, at least in part, motivated by the features of each type of resource. For the most part, for example, conventional cinematic resources are almost always non-interactive and screen-based (i.e., they require a screen for display purposes). Unconventional cinematic resources, by contrast, are almost always interactive and are much less dependent on a display screen.²

Relative to this taxonomy, *Blade Runner 2049* is cast as a conventional cinematic resource (i.e., an instance of a movie), while *Joi* is cast as a

(fictional) instance of a hologram and thus a type of unconventional cinematic resource. But now notice something important. In addition to her status as an unconventional cinematic resource, Joi is also one of the characters in *Blade Runner 2049*. In this sense, she qualifies as a cinematic entity (or cinematic object)—a cinematic resource that is depicted in the context of another cinematic resource.

All of this raises an issue about the scope of cinematic philosophy. If cinematic philosophy is concerned with the philosophical significance and impact of the cinematic medium, and holograms count as cinematic resources, then holograms (and other unconventional cinematic resources) ought to fall within the scope of cinematic philosophy.

There may, however, be some resistance to this idea. Perhaps, for example, it is unclear whether holograms have any philosophical significance, or, if they do, whether that significance is germane to cinematic philosophy. The beauty of *Blade Runner 2049*, I suggest, is that it helps to resolve this issue in favour of the idea that holograms ought to form part of the subject matter of cinematic philosophy.

To help us creep up on this idea, let us identify two views about the scope of cinematic philosophy. The first of these is what we might call the *traditionalist view*. This view insists that cinematic philosophy ought to limit its attention to the realm of conventional cinematic resources. The second view is what we might call the *expansionist view*. This view advocates an expansion in the traditionalist view to accommodate the class of unconventional cinematic resources.

Of course, in the context of *Blade Runner 2049* the distinction between traditionalists and expansionists is not particularly crucial. This is because Joi is a legitimate focus of interest for the proponents of both views. Given that Joi forms part of a conventional cinematic resource (a movie), she will be of mutual interest to the proponents of either position, since neither the traditionalist nor the expansionist disputes the relevance of conventional cinematic resources to cinematic philosophy. At the same time, however, Joi is not the sort of cinematic resource that a fan of the traditionalist position would find particularly interesting once we move beyond the borders of the big screen. For the most part, only the proponent of the expansionist view will see real-world hologrammatic resources as falling within the intellectual orbit of cinematic philosophy.

At this point, however, the traditionalist position starts to look a little problematic. For, aside from the fact that movies are projected on the

big screen and holograms are projected in peri-personal space (possibly without a screen!), it is not particularly clear what it is that gives conventional cinematic resources their seemingly distinctive philosophical significance. This is especially so when the focus of our philosophical attention is directed to a movie that reveals the philosophical significance of a somewhat different type of cinematic resource. By depicting a mixed-reality environment, for example, *Blade Runner 2049* encourages us to consider the philosophical significance of mixed-reality resources, and it does so irrespective of whether or not we commit ourselves to the traditionalist or expansionist camp. But once we have identified these points of philosophical significance, our attention can then turn to the real-world counterparts of the cinematic resources depicted in the movie. This is surely just as important as the analysis of those resources in a big-screen context. In fact, the things that make Joi interesting from a philosophical standpoint are, I suggest, just as applicable to an emerging array of mixed-reality technologies. This is important, for it helps us identify a point of philosophical significance about *Blade Runner 2049*. By serving as the legitimate target of philosophical analysis for both traditionalists and expansionists, Joi helps to lay the foundation for a broader (meta)philosophical debate about the future scope and direction of cinematic philosophy. She does this by speaking to the expansionist agenda through the lens of a traditionalist analysis. In particular, it is via a traditional form of cinematic philosophical analysis (such as the one you are now reading) that the possibility and potential of an expanded cinematic philosophy starts to come into sharper focus.

Some of the philosophical issues raised by the study of unconventional (especially hologrammatic) cinematic resources are discussed in the next section. For now, however, it is worth noting the way in which Joi's status as a hologrammatic entity resonates with our current technological, commercial, and scientific interest in what is sometimes called *X Reality* (XR) (also known as *cross reality*)—a catch-all term for the various forms of “reality” picked out by the notions of cinematic, augmented, mixed, blended, and virtual reality. Indeed, with the advent of holographic devices, such as Microsoft's HoloLens (see Figure 7.3), Joi appears less like a beguiling fantasy advert suspended above the rain-drenched streets of a fictional cinematic universe, and more like a poster child for the next generation of personal assistants, digital companions, and (Who knows?) virtual lovers. In this respect, the cinematic reality



Figure 7.3 Microsoft HoloLens promotion image.

of *Blade Runner 2049* raises important questions about the effect of unconventional cinematic resources on the shape of our own physical and social realities. If this is, at least in part, one of the things that makes *Blade Runner 2049* interesting from a philosophical perspective, then why demur from the proposal that hologrammatic resources ought to form part of the theoretical and empirical agenda of (an expanded) cinematic philosophy?

At this point, it is probably up to the proponent of the traditionalist view to specify why they regard conventional cinematic resources as worthy of specialist (and exclusive) philosophical treatment. For a variety of reasons, I very much doubt that any sort of defence of the traditionalist position can be made to work. And this remains the case, even if we resort to alternative ways of labelling the relevant philosophical enterprise (e.g. the *philosophy of film*). To help us understand why this relabelling strategy is unlikely to work, consider that *Blade Runner 2049* depicts a mixed-reality world, and it raises philosophical issues about a particular kind of mixed-reality resource within that world, namely, holograms. Thus, even if we limit our attention to the big screen, there are still times when the reach of cinematic philosophy will extend to the realm of mixed-reality resources, and it is unclear why the big-screen context ought to serve as the ultimate arbiter of philosophical interest. To help reinforce this particular point, it is perhaps worth noting the ease

with which *Blade Runner 2049* can itself be rendered as a hologrammatic resource. There is no reason, for example, why *Blade Runner 2049* cannot be viewed using the built-in augmented reality video player that ships with Microsoft HoloLens. At this point, the erstwhile crisp distinction between the realms of conventional and unconventional cinematic resources may itself start to look a little blurry (or ambiguous?). For what are we watching when *Blade Runner 2049* is rendered as a hologrammatic resource in peri-personal space? Is it a movie, a hologram, a hologrammatic movie, or all of the above?

It might be said, of course, that a distinctive feature of conventional cinematic resources (and the feature that marks them out for specialist philosophical treatment) is that they have a rich narrational structure. But rich narrational structures are not a universal feature of the movies that have been targeted by cinematic philosophy (see Wartenberg, 2007, chap. 7), and, in any case, there is no reason why unconventional cinematic resources are ill-equipped to feature narrational elements. Video games, such as *Red Dead Redemption 2*, for example, allow human game-players to create their own narrative structures, courtesy of their capacity to exert “directorial” control over events within a virtual world. And if Joi is akin to a character in a computer game, then isn’t K doing something similar to a human game-player—creating a narrative, courtesy of his interactions and exchanges with a virtual character? The only substantive difference here is that Joi is part of K’s life story—part of his personal narrative history and part of his narrative arc in the movie. Given that such forms of narratological entanglement are apt to be a source of philosophical interest in the context of *Blade Runner 2049*, it is unclear why similar forms of narratological entanglement should fail to be a source of philosophical interest just because they involve resources that lie beyond the movie theatre walls.

Minds, mates, and the moving image

JOI: An emanator.

JOI: Thank you.

K: Honey, you can go anywhere in the world you want now.
Where you want to go first?

K's question is particularly apt, for having broached the idea of an expanded cinematic philosophy, I am now faced with the task of shedding some light on this newly unveiled region of the philosophical terrain. The primary problem here is not so much one of scarcity as it is superabundance. For there are lots of things to see (and do) in this particular neck of the philosophical woods and, unfortunately, the tour bus can only visit so many locations.

I will limit my attention to three issues. These were selected according to a number of criteria. First, the issues speak to a general concern with the cognitive value of cinematic resources. In particular, they begin to show, at least in outline form, how an expanded set of cinematic resources (most notably, hologrammatic resources) might influence the shape of human cognitive processing. Second, the issues are, in one way or another, inspired by my own experience of watching *Blade Runner 2049*. Third, the issues begin to illuminate the nature of the broader interdisciplinary links that connect cinematic philosophy with work in a number of other disciplines, such as cognitive science, computer science, and the philosophy of mind.

- (1) **Cinematic AI:** Joi is both a cinematic entity (a hologram) and a form of AI. This raises questions about the extent to which intelligent systems might be implemented in the cinematic medium.
- (2) **Virtual Companions:** Contrary to the traditional image of AI systems as possessing superhuman cognitive abilities, Joi's intelligence is firmly rooted in the social domain. She is intended to be a companion for K—a source of emotional comfort and socio-sexual stimulation. This raises questions about the extent to which socially adroit hologrammatic AI entities ought to be seen as virtual additions to the human social world.
- (3) **Hologrammatically Extended Minds:** I suggest that holograms provide the basis for a particular form of cognitive extension (Clark, 2008; Clark & Chalmers, 1998), which involves the integration of (unconventional) cinematic resources into cognitively relevant processing loops. The resulting cognitive organisations are what we might call *hologrammatically extended minds*. Hologrammatically extended minds are, in essence, cognitive systems that are constituted, at least in part, by resources that fall within the scope of (an expanded) cinematic philosophy.

Let us take these points in turn. The first issue relates to the notion of cinematic AI. We have already seen why Joi (and thus her real-world hologrammatic kin) ought to be regarded as a particular kind of cinematic resource (see above). There should also be little doubt that Joi qualifies as a form of AI system. She is, after all, a technological artefact, produced by Wallace Corp., and she is clearly capable of behaving in an intelligent manner. The result is that Joi qualifies as a form of cinematic AI. She is, in short, emblematic of the idea that certain kinds of cinematic resource (e.g. holograms) are able to function as AI systems.

Despite the seeming banality of this claim, it is, in fact, of crucial importance. In particular, it suggests that the cinematic medium is not just a medium in which philosophically useful work might be done; it is also a medium that supports the realisation of (sentient?) intelligent systems. The resulting set of issues are ones that not only straddle the ostensibly distinct intellectual terrains of cinematic philosophy, cognitive science, and the nascent field of holographic computing, they are also ones that might be seen to blur the very boundaries that separate these particular fields of intellectual enquiry.

Joi's status as a virtual companion also serves as a point of philosophical interest. In particular, Joi is a system whose intelligence is primarily oriented to the socio-emotional domain. She is a system designed to fulfil a social function: a virtual companion that aims to act as a technological filler for whatever social, emotional, and sexual gaps exist in the lives of those who purchase her.

There is, of course, no real-world parallel to Joi. Virtual assistants, such as Amazon's Alexa, Apple's Siri, and Microsoft's Cortana, are testament to the progress that has been made in speech recognition, conversational interaction, and information retrieval. But, for the present at least, these systems remain little more than voice-based intermediaries to an online world of digital information and networked devices.

This doesn't mean that virtual companions are forever consigned to the realms of fantasy and fiction, however. As is evidenced by work into social (and sexual) robots, believable game characters, and digital companions, there is no shortage of scientific (and commercial) interest in the possibility of socially adroit AI systems. Assuming that research in these areas merges with that in holographic computing, it is not inconceivable that Joi-like systems might one day feature as elements of human social reality. This is important, for we typically think of holograms

(*qua* mixed-reality resources) as a means of adding information to the local environment and thereby altering the nature of our cognitive and experiential contact with physical reality. The beauty of *Blade Runner 2049* is that it directs our attention to the social impact (and relevance) of hologrammatic resources—the way in which some types of XR resource may be poised to alter (extend? enrich? augment? contaminate?) the structure and dynamics of human social life.

The third and final issue on our list concerns the notion of hologrammatically extended minds. Hologrammatically extended minds arise as the result of the incorporation of a hologrammatic resource into a materially extended cognitive circuit. A hologrammatically extended mind is thus a particular kind of extended cognitive system (Clark, 2008) or extended mind (Clark & Chalmers, 1998). The only difference between a hologrammatically extended mind and a conventional extended mind is the nature of the extra-organismic resource that is factored into a cognitively relevant information-processing loop. In the case of hologrammatically extended minds, these resources are a particular kind of cinematic resource, namely, a hologram. (Holograms are, of course, not the only kind of cinematic resource that may be incorporated into a materially extended cognitive circuit; nevertheless, they are the most important kind of cinematic resource relative to our interests in the present chapter.)

Although extended minds have been the focus of a long-standing (and ongoing) philosophical debate, the specific notion of a hologrammatically extended mind is not one that seems to have garnered much philosophical attention. This is unfortunate, for hologrammatically extended minds yield a number of puzzles and problems that extend the traditional palette of philosophical concerns in this area. Note, for example, that issues of cognitive extension are typically discussed with reference to concrete, physical resources, such as notebooks and iPhones (e.g. Clark & Chalmers, 1998). In this case, the philosophical polemic centres on the nature of the functional relationship that exists between resources that lie either side of the organismic boundary. The question is whether we should talk of some extra-organismic resource (such as a notebook) as forming part of the mechanistically relevant fabric that realises some cognitive state or process. This issue remains important in the case of hologrammatically extended minds, but we now have an additional concern: To what extent does it make sense to talk of a virtual cinematic

resource as forming part of the physical machinery of the mind? One issue here relates to the “virtual” status of a cinematic resource. How can a photic resource (i.e., a resource made of light) feature as part of the causally active physical fabric that realises cognitive states and processes? Does this mean that hologrammatically extended minds are constituted, at least in part, by light? Is it possible for cognitive routines to emerge from the forms of photic flux that define the moving image? And, if so, what does this mean for our current understanding of the cognitive and philosophical value of the cinematic medium?

There are, no doubt, many issues that are raised by the possibility of hologrammatically extended minds, and I will not attempt to rehearse those issues here, let alone proffer anything in the way of a solution. For present purposes, what matters is merely the fact that we have identified a point of convergence between cinematic philosophy and the philosophy of mind, one that expands the nature of the debates in both areas, while simultaneously drawing inspiration from the cinematic spectacle that is *Blade Runner 2049*. Having said that, there are no specific examples of hologrammatically extended minds that I can point to in *Blade Runner 2049*. The functional contribution of *Blade Runner 2049*, in this respect, is not so much to illustrate a specific philosophical claim (see Wartenberg, 2009) as it is to inspire a new way of thinking about the philosophical significance of the cinematic medium. There is, no doubt, an important sense in which my own biological brain did much of the heavy lifting in order to bring this particular idea to light, and yet it was the play of light and sound (on the big screen) that managed to coax my biological brain into doing some (hopefully) useful philosophical work. (No mean feat, given the typically torpid nature of my neural machinery.)

This does not mean that the notion of hologrammatically extended minds is irrelevant to our philosophical assessment of *Blade Runner 2049*. For, once we have the idea to hand, we can begin to put it to philosophical work in assessing some of the issues raised by the movie. It might be thought, for example, that nothing like Joi could exist outside the realms of *Blade Runner*’s fictional reality: “How could a mere hologrammatic projection count as anything like a cognitive system, let alone a virtual person!?” From a mechanistic perspective, however, it is arguably a mistake to view Joi solely as a holographic projection. For Joi’s projection is rendered by a physical device (e.g. the emanator), and this should arguably be included as part of the material mix that makes Joi the peculiar

cognitive entity she is. Joi is, in this sense, a hybrid entity—a mix of the physical and the virtual: part physical computation, part cinematic projection. Her claim to cognitive status is thus no more (and probably no less) contentious than is the basic possibility of a hologrammatically extended mind. For, if the material fabric that undergirds our own (experientially charged!) forms of cognising may, on occasion, include a mix of both the physical (the biological) and the virtual (the cinematic), then why object to the idea that Joi's own cognitive wherewithal (sentient or otherwise) may, on occasion, rely on a material fabric that features both a physical device and that most mesmerizing of human cultural innovations—the play of light that is the moving image?

Notes

- 1 At the beginning of the movie, Joi says (to K) “I’m so happy when I’m with you.” It is only later in the movie, immediately prior to her “death,” that Joi attempts to proclaim her love for K.
- 2 Volumetric holograms, for example, do not require a display screen, nor do resources that are projected using Virtual Retinal Display (VRD) devices.

References

- Clark, A. (2008). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. New York: Oxford University Press.
- Clark, A., & Chalmers, D. (1998). The Extended Mind. *Analysis*, 58(1), 7–19.
- Gaut, B. (2010). *A Philosophy of Cinematic Art*. Cambridge: Cambridge University Press.
- Hyslop, A. (2018). Other Minds. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2018 ed.). Stanford: Stanford University.
- Nagel, T. (1974). What Is It Like to be a bat? *The Philosophical Review* 83(4), 435–450.
- Shamir, T. S. (2016). *Cinematic Philosophy*. Basel: Palgrave Macmillan.
- Smalley, D. E., Nygaard, E., Squire, K., Van Wagoner, J., Rasmussen, J., Gneiting, S., et al. (2018). A Photophoretic-Trap Volumetric Display. *Nature* 553(7689), 486–490.
- Tormey, J. F., & Tormey, A. (1983). Art and Ambiguity. *Leonardo* 16(3), 183–187.
- Wartenberg, T. E. (2007). *Thinking on Screen: Film as Philosophy*. Abingdon: Routledge.
- Wartenberg, T. E. (2009). Film as Philosophy. In P. Livingstone & C. Plantinga (Eds.), *The Routledge Companion to Philosophy and Film* (pp. 549–559). Abingdon: Routledge.