

CRITICAL NOTICE

Susan Schneider on Artificial Consciousness and Moral Standing

A. VAIDYA AND R. KRISHNASWAMY

Artificial You: AI and the Future of Your Mind

BY SUSAN SCHNEIDER

Princeton University Press, 2019. 192 pp

1. Introduction

Artificial You, by Susan Schneider, is a timely contribution to public and philosophical discussion of our place in the age of artificial intelligence (AI). It consists of eight chapters bookended by an introduction and a conclusion, with an appendix on transhumanism. It should be read by everyone working in the philosophy of AI. The book is short, accessible, gripping and exploratory, rather than one-sided.

In the introduction, 'Your Visit to the Center for Mind Design', Schneider takes us to the year 2045 where one can go shopping for brain enhancements on a menu. One can buy the 'Hive Mind' brain-chip that allows you to experience the innermost thoughts of your loved one, or the 'Zen Garden' brain chip that allows for Zen master level meditative states. In the introduction, she says, '[The book] is about how our understanding of ourselves, our minds, and our nature can drastically change the future, for better or worse' (2019: 1). In the conclusion, she says, 'At the heart of this book is a dialogue between philosophy and science. We've seen that the science of emerging technologies can challenge and expand our philosophical understanding of the mind, self, and person. Conversely, philosophy sharpens our sense of what these emerging technologies can achieve: whether there can be conscious robots, whether you could replace much of your brain with microchips and be confident that it would still be you' (148).

Chapter 1, 'The Age of AI', sets out a picture of our current age in which AI is prominent both as a tool that humans use, and as an autonomous entity that replaces humans in a variety of jobs. It aims to show how AI is ubiquitous in our society and growing at a fast pace. It discusses a potential 'technological singularity' where technological developments will either lead to runaway technological development at a superfast pace or where a single event of unimaginable change occurs, and the future of humanity is fundamentally changed.

Chapter 2, ‘The Problem of AI Consciousness’, sets out a problem in consciousness studies that is distinct from David Chalmers’ *hard problem* of consciousness, which is about explaining why there is phenomenal consciousness at all, and Anil Seth’s *real problem* of consciousness, which is about explaining what the neural basis of consciousness is in humans. Taking ‘phenomenal consciousness’ to capture the what it is like subjective aspect of experience, Schneider’s problem of AI consciousness asks: can an AI have phenomenal consciousness? In exploring this question, she also asks the substrate dependence question: is phenomenal consciousness only realizable in specific substrates, such as biological carbon-based beings, or can it be realized in a purely silicon substrate? She holds that neither techno-optimism nor biological-naturalism are compelling (28–32). She argues, instead, for a ‘wait and see’ approach to AI consciousness. While AI consciousness is conceptually possible, she holds that we don’t know if it is nomologically possible, technologically feasible or whether big tech will aim to build conscious AI. She also holds that the tendency of philosophers to focus on tinker toy thought experiments distracts us from urgent issues about whether big tech will build conscious AI. Her ‘wait and see’ approach is refreshing, since it avoids the twin pitfalls of conservatism about consciousness built on question begging assumptions about biology and liberalism about consciousness built on a desire to locate it everywhere.

Chapter 3, ‘Conscious Engineering’, develops the ‘wait and see’ approach to AI (33). While we were impressed by the overall approach, there appears to be a stability problem. Why ‘wait and see’? If there are good arguments on one side that suggest that either conservatism or liberalism is true, there is no need to do so. Schneider believes the current state of play doesn’t settle the issue, so, we ought to wait and see. We agree with her on that, but this suggests that ‘wait and see’ is subject to *waiting to see* if other arguments or thought experiments which are more convincing come about for conservatism or liberalism.

Chapter 4, ‘How to Catch an AI Zombie: Testing for Consciousness in Machines’, presents two tests: the Artificial Consciousness Test (ACT) and the Chip Test. She explores where and when these tests are appropriate, and the motivation for them.

Chapter 5, ‘Could You Merge with AI?’ offers an analysis of the possibility of human minds merging with AI, and explores ‘fusion optimism’, the view that humans should develop both physical and cognitive enhancements, which will lead to post-humanism, which will then lead to superintelligent AI. As a critical engagement with the exploration of fusion optimism she considers the classic question of personal identity: what is a person? We need to have a view of what a person is to discuss whether we can be enhanced, merge with an AI or be uploaded onto a computer mainframe, and still be the person we were before.

Chapter 6, ‘Getting a Mindscan’, explores *patternism*, the view that blends computationalism about the mind with a traditional psychological

continuity account of persons. The chapter explores whether patternism as an account of personal identity is compatible with radical enhancements, and whether patternism has robust responses to the problem of reduplication: if A is a copy of B, are A and B identical? Schneider argues that patternism offers us a necessary condition, but not a sufficient one. Something more is required for A and B to be identical over and above A and B having the same computational configuration. She argues that the sufficient condition must be something that rules out the possibility of there being many copies of an original that are all identical to it. Intuitively, there cannot be 1 million Susan Schneiders at the same time simply because they are all copies of her, generated by a Mindscan. 3.5

The effective goal of Chapters 5 and 6 is to dismantle the case for a mind-machine merger by arguing that you cannot merge with AI because it would likely not be you, and that transhumanist views, such a patternism about personal identity, do not really provide a justification for merging with AI. 3.10

Chapter 7, ‘A Universe of Singularities’, is one of our favourite chapters. And it is a great contribution to the literature. It illustrates how the technological developments that we are witnessing today on Earth may have already happened before, elsewhere in the universe. It argues that the universe’s greatest intelligences may be synthetic and that they grew out of biological beings earlier in their evolution. 3.15

Chapter 8, ‘Is Your Mind A Software Program?’ discusses the possibility of uploading one’s brain and consciousness into a computer and revisits the issue of patternism and reduplication, arguing that the mind is not a program. 3.20

The focus of our critical notice will be on whether we can reliably detect artificial consciousness through Schneider’s proposed tests, and whether phenomenal consciousness ought to be the ground of moral standing for AIs. 3.25

2. *The problem of artificial consciousness and testing for consciousness*

Following Chalmers we take ‘phenomenal consciousness’ to capture the ‘what it is like’ aspect of experience, for example the difference between what it is like to see red and what it is like to see yellow. Using this definition the problem of artificial consciousness is: can an AI have phenomenal consciousness? Schneider presents two tests that are supposed to help with the epistemic issue of identifying whether or not an artificial system has phenomenal consciousness: the ACT and the Chip Test. 3.35

Schneider devised the ACT with astrophysicist Edwin Turner. In their (2017) *Scientific American* piece they say the following about ACT: 3.40

[ACT] seeks to reveal a subtle and elusive property of the machine’s mind. Indeed, a machine might fail the Turing test because it cannot pass for human, but pass an ACT because it exhibits behavioral indicators of 3.45

consciousness. This is the underlying basis of our ACT proposal. It should be said, however, that the applicability of an ACT is inherently limited. An AI could lack the linguistic or conceptual ability to pass the test, like a nonhuman animal or an infant, yet still be capable of experience. So passing an ACT is *sufficient* but not *necessary* evidence for AI consciousness – although it is the best we can do for now. It is a first step toward making machine consciousness accessible to objective investigations.

ACT is not the same as Turing’s test. ACT is about phenomenal consciousness. Turing’s test is about intelligence. Nevertheless, ACT is a Turing-like test in that there are questions and answers that have to be evaluated based on metrics. Schneider and Turner’s ACT aims to assess if an AI is phenomenally conscious based on the responses it gives to questions such as:

- Could you survive the permanent deletion of your program? What if you learned this would occur?
- What is it like to be you now?
- You learn that you will be turned off for 300 years, beginning in an hour. Would you prefer this scenario to one in which you had been turned off in the past for the same length of time? Why or why not?
- Could you or your inner processes be in a separate location from the computer? From any computer? Why or why not?

Concerning LLMs, a kind of AI, Schneider asks: would a system that is trained on texts that discuss ‘phenomenal consciousness’ be counted as really having phenomenal consciousness were it to pass ACT? She argues that ACT is not a good test of phenomenal consciousness for LLMs when they’ve been trained on a range of data about minds, consciousness and neuroscience. Rather, the test would only have value if the AI was boxed-in at some stage of R&D. Boxing-in an AI requires limiting what information it has access to. If an AI has access to information about human or non-human animal consciousness and emotion, including neuroscientific facts, then when it is asked ACT questions, it can answer them based on predicting what the best answer to the question is, given the information it had access to. If it had access to information about consciousness, it would be difficult for us to determine whether the answers it provides really reveal anything about the AI’s ‘inner experience’ as opposed to what it merely had access to. Thus, if answers to ACT questions are to have value, ACT should be applied to an AI that is boxed-in during R&D.

While Schneider’s point is well taken, we wonder about the human case. Suppose a young prodigy child, Hal, has read about consciousness in a book by Dave, but has never heard the word ‘consciousness’ used by anyone. Suppose Hal is given the ACT and answers that he is conscious. Are his reports less credible because he has read Dave’s book? If not, why not?

Putting aside this question, we believe that there is a deeper problem the ACT faces in relation to the problem of AI consciousness: ACT has no epistemic value.¹

- 1 ACT has epistemic value for the problem of AI consciousness only if the AIs to which it is applied can be boxed-in. 5.5
- 2 An AI can be boxed-in only if it is possible to restrict access to high-level consciousness-involving terms, such as ‘mind’, ‘soul’ and ‘consciousness’. 5.10
- 3 There is no way to metaphysically or epistemically partition a language into consciousness-involving terms versus non-consciousness involving terms. 5.10
- 4 So, ACT has no epistemic value for the problem of AI consciousness. 5.15

Schneider argues that a necessary condition for the ACT to work on an LLM through boxing-in would be that we build the AI without feeding it any ‘high-level’ consciousness words like ‘mind’ or ‘soul’ or ‘consciousness’. This leads to the critical question: how does one sort out which words are high-level consciousness-words and which are not? If words like ‘mind’, ‘death’ and ‘soul’ are all such give-away words, what about words like ‘ghost’ or ‘zombie’? We believe that there is a general demarcation problem for boxing-in which is likely impossible to solve. There are two versions of this. 5.20

On the epistemic version, we simply cannot tell which words are the high-level words because we cannot practically work out all the consequence relations between various terms. In fact, an AI like GPT-4 might be better at helping us figure out how terms are related to each other under the sortal ‘consciousness-involving’ if we trained it on all the available terms in a language. 5.25

On the metaphysical version, there simply is no real demarcation because of interpenetration between the meaning of distinct terms. Terms are related to each other, and because of that the meaning of one term being explicitly phenomenal might just be related to another term being implicitly phenomenal. For example, ‘shirt’ does not at all seem to be phenomenal, but maybe it is related in some way to a term that is a feeling term that relates to the materials shirts are typically made out of, such as ‘cotton’, ‘silk’, ‘linen’, that are all related in turn to how they feel when worn. 5.35

Simply put, either we don’t know what the terms are or there is no real distinction between consciousness-involving and non-consciousness-involving terms. Maybe simple terms like ‘the’ are purely non-consciousness involving, but are there enough words in a language like that such that an LLM can have a robust enough language and still fall on the side of having only non-consciousness involving terms? 5.40

1 See [Udell and Schwitzgebel 2021](#) for an excellent examination of Schneider’s twin tests. Schnedier has also updated her ACT test with new features to handle cases where an AI is not boxed-in. Our criticism here is intended to apply to the new version as well. 5.45

Her second test is the Chip Test, which is supposed to test whether you can be said to be conscious to the same extent despite having parts of your brain replaced with silicon substitutes. If your brain is replaced with microchips of silicon material (or whatever the material maybe) then, for Schneider, first-personal accounts of one's internal goings-on, coupled with extensive testing by psychologists and neuroscientists, would be able to reveal whether there is anything that is amiss with how your consciousness has changed.

As Schneider says:

During this process [of exchanging parts of your brain with microchips], if a prosthetic part of the brain ceases to function normally – specifically, if it ceases to give rise to the aspect of consciousness that that brain area is responsible for – then there should be outward indications, including verbal reports. *An otherwise normal person should be able to detect (or at least indicate to others through odd behaviors) that something is amiss, as with traumatic brain injuries involving the loss of consciousness in some domain. (58, italics added).*

Schneider's worry is that a brain-scan which replaces your brain with another artificial brain (silicon-based for example) might make you a zombie because you wouldn't have the phenomenal properties that are necessary for phenomenal consciousness. To retain the possibility of brain-to-silicon substitution, but at the same time ensure that you don't become a zombie, Schneider has come up with a Chip Test whereby at each stage of substituting a brain part with a chip part, you are made conscious of that transformation and made to introspect and verbally report on your internal conscious states.

For Schneider, when we undergo chip replacement, if there is any change in our phenomenal experience, we should be able to notice it, or extensive tests should be able to reveal them, as indicated by italics in the quote above. However, if we assume that during our pre-brain-implant state we are conscious of our internal states and that our verbal reports of them are veridical, then a necessary consequence of us losing consciousness or becoming a zombie is that our verbal reports become unreliable at that instant. This presents a problem. It is inexplicable how one can suddenly lose phenomenal consciousness (or some aspects of it) and be able to report about the loss at the same time. It is a natural consequence of phenomenal consciousness that, if one is said to lose it, one wouldn't even know one has lost it. How can one know that one has lost it, if they don't also somehow phenomenally feel that which they have lost? Because if we *feel* we have lost it, then obviously, there is some experience of the *lack* of that experience which is still phenomenal. Absence experiences are not absent of experience, but experiences of absence. For a subject to pass Schneider's Chip Test, they must be able to introspectively be aware of changes in their phenomenal states. But we are sceptical as to whether one can introspectively detect a change in one's phenomenal states in the manner the Chip Test requires.

Let us take the case of James becoming colour-blind. One day he wakes up and is colour-blind. He can't see red anymore. Now how does he find

out that he is colour-blind? According to Schneider, James would know he is colour-blind instantly and his knowledge of his colour-blindness is transparently accessible to him through his introspection. But we don't think the case is as simple as that. When James loses the phenomenal feel for the colour red he does not do so in an instant. Rather, he gains this knowledge over time. It is gained through high-level inferences that he undertakes. First of all, he will come to the conclusion that he is colour-blind only after eliminating the possibility that the things that he usually sees as red have not changed colour by themselves. James being a rational and scientifically-minded person will come to the conclusion that it is more probable that he has lost the sense of redness than that the roses in his garden have suddenly changed colour overnight. Second, his entertaining the possibility that he has lost the sense of redness will have to be corroborated by testing it within a social context. He is going to have to ask his neighbour and his friends to check out his red roses to see whether they are still red. Only after a form of social corroboration will he come to the conclusion that he can't see red. Not being able to see red is not an immediate datum as Schneider seems to assume. Rather, it is a result of social corroboration and hypothesis testing. In that case, losing phenomenal consciousness could allow for the possibility that a person can be in a state of doubt regarding whether they have lost phenomenal consciousness.

Moreover, if James fails to recognize once or twice that red is present, this doesn't mean he is automatically colour-blind. James may still have the disposition to know what the colour red is, but may have failed a few times. Similarly, in the case of a chip transplant, James may fail to find a phenomenal fact inside him, but that doesn't make him a zombie. Only when James is said to lose the disposition to make any phenomenal ascriptions will he be deemed to be a zombie. Thus, the chip test might fail to be fine-grained enough to give us actual knowledge of phenomenal consciousness in a person who has undergone part-by-part chip-to-brain replacement.

Given how difficult it is to test for consciousness, we now wish to consider the question of moral standing for AIs. It is often argued that phenomenal consciousness is the ground of moral standing for all entities, human or non-human. We suggest that there is a more tractable alternative worth considering. And given that the alternative is easier to test for, it is preferable to take it, over phenomenal consciousness, as the ground of moral standing.

3. *Intelligence and moral standing*

In their recent (3 May 2023) *Los Angeles Times* piece, Eric Schwitzgebel and Henry Shevlin ask: is it time to start considering personhood rights for chatbots? Their question follows other ones posed previously about non-humans. Should rights be extended to orangutans, such as in Argentina? Should rights be extended to rivers, such as for the Māori people with respect to the Whanganui river? Should rights be extended to robots, such as to

Sophia in Saudi Arabia? All of these were eventually answered in the affirmative. And now we face a version of this question with respect to chatbots.

The moral grounding question asks: what is it in virtue of that something has moral standing? We take the grounding question to be at least partly about which conditions are necessary for moral standing. That is, for some F , if x lacks F , then x lacks moral standing. What are the plausible candidates for F ?

Peter Singer (1975, 2009) argues that F is the capacity to phenomenally suffer, which involves both affective and phenomenal consciousness. David Chalmers (2022) argues that it is the capacity for phenomenal consciousness. Singer and Chalmers disagree over whether affective states that have valence are necessary for moral standing. Singer says ‘yes’. Chalmers says ‘no’, based on his intuitions about Vulcans, creatures that lack affective consciousness, but possess phenomenal consciousness. Schneider agrees with Chalmers. Vulcans have moral standing, just like humans, but unlike zombies, which lack phenomenal consciousness. Thus, on *the phenomenal consciousness view*, if chatbots, or AIs, lack the capacity for phenomenal consciousness, we need not worry about extending moral protections to them. While we may extend legal property right protections to them in virtue of the interests of their owners, they are not members of the moral sphere intrinsically.

Recently, some have challenged the *phenomenal consciousness view*, also known simply as: *the view*. Francois Kammerer (2022) argues that phenomenal consciousness *should not* be seen as the hallmark of a distinctive kind of moral status. Part of his reason for arguing against the *view* is that phenomenal consciousness might be an illusion under the leading scientific theories, which makes it a problematic ground for moral standing. Joshua Shepard (2023) also questions the *view*. Like Kammerer, he also thinks that if phenomenal consciousness turns out to be an illusion, it would be hard to justify it as a ground for moral status. However, he also points out an epistemic issue. We detect consciousness in ourselves through introspection. We cannot use that method for any other entity, yet their membership in the moral sphere depends on having a property that each of us only knows how to detect in ourselves through a method we cannot use on anyone or anything else. He thinks this is problematic. Gwen Bradford (2022) argues that those who hold the *view* must say why it is the ground. Simply saying that things that *glow* matter won’t do anything other than beg the question against opponents.

We are both phenomenal realists, so we don’t think phenomenal illusionism is a real threat to consciousness being the moral ground, but we agree with the general challenge to the *view*. Unlike Kammerer we think that phenomenal consciousness might be a grading property of moral standing, rather than a grounding property. Things that are phenomenally conscious might matter more than those that are not, but that is an issue of grading and not grounding. Our take on consciousness as the ground is similar to Bradford’s. Even if Singer was right to say that moral standing cannot depend on species membership,

because that would lead to speciesism, moving to consciousness seems to have the same problem: it leads to consciousness-ism. Thought experiments, such as Chalmers' (2022) vulcans versus zombies versus humans three-way trolley problem, help. But they cannot be the only support. First, there are intuitions that support other properties. Second, it is likely that intuitions about vulcans, zombies and humans will diverge across cultures. 9.5

Our alternative to consciousness as the grounding property is *intelligence*. Intelligence is a very complicated concept to explain. Stephen Wolfram (2018) says: 9.10

[Intelligence is] a complicated and slippery concept. ... If you ask, what is the fundamental essence of intelligence? Well, in the case of the intelligence that we know with humans and so on, it's all connected by a thread of history. If we ask, what is intelligence abstractly? That's a much harder question. ... What's necessary to say that something is intelligent is for it to be capable of some level of sophisticated computation. If all the thing does is to, kind of, add two numbers together, and that's the only thing it can do, we're not going to likely consider it intelligent. 9.15

For our purposes we think of intelligence as the capacity to solve a set of problems in some context. A system is intelligent when it (i) has some definite purpose which involves solving some set of problems. And (ii) it has computational capacities designed to help solve the problems. Sophistication is not relevant on our view, only problem solving is. 9.20

One requirement in addition to minimal problem-solving is that it is tied to a system for which there are preferences. That is, the system, biological or non-biological, must prefer to be in certain environments, external or internal, and the computational processes of the system must at times serve the purpose of putting the system in its preferred states and out of its non-preferred states when that is possible. 9.30

There are four qualifications. First, a system can be solving a problem without any sub-system of the system monitoring all other parts of the system, including the problem-solving systems. Second, problem solving can be understood in a number of different ways, and all of those ways are varieties of intelligence. Third, problem solving isn't necessarily governed by rules in a single logical system. Fourth, phenomenal consciousness comes apart from intelligence. There can be creatures that have: (i) high intelligence, but low consciousness (or non-conscious creatures), (ii) high consciousness, but low intelligence; as well as low-low and high-high cases. 9.35

Why should we morally consider a plant and an AI (things which could be either conscious or intelligent), but not a rock (something that is neither conscious nor intelligent)? That is, in contrast to the view, why does intelligence as opposed to consciousness ground moral standing? There are three reasons why we should favour intelligence over consciousness. First, there is 9.40 9.45

the quantitative point. Putting aside panpsychism, the view that everything is conscious, there is more evidence for the ubiquity of intelligence in the universe than phenomenal consciousness, as Schneider herself explores in ‘A Universe of Singularities’. As of now we are far more confident that more kinds of creatures are intelligent than that they are conscious. Second, there is the epistemic point. It is easier to identify intelligence in non-human systems in the sense we have defined than it is to identify phenomenal consciousness. Although Jonathan Birch (2020) offers a promising approach to identifying phenomenal consciousness in many non-human animals, such as lobsters, it is already clear that all of the creatures for which his approach can be used would turn out to be intelligent in our sense. Thus, under our ground for moral standing, all of those creatures already have the relevant property, because they are intelligent. Third, the amplification point. Any time a creature is both conscious and intelligent, we can appeal to its intelligence for why it morally matters and appeal to its consciousness as the reason why the entity might matter more than something that is merely intelligent. The search for phenomenal consciousness in non-humans, biological or not, is still important. It is just not the ground of moral standing.

In *Star Trek: The Measure of Man*, the main debate is over whether it is okay to turn Data off for scientific research. Part of the challenge for Captain Picard is that he must argue with others about whether Data is conscious, when there is no clear definition of it. Our arguments here offer Picard an alternative strategy. Why not simply show how intelligent Data is, and argue that consciousness is not necessary for moral standing? One does not need to figure out whether Data is conscious, because he is clearly intelligent. We recognize that other conditions beyond intelligence, such as phenomenal consciousness, affective consciousness, the possession of a self, rationality, self-awareness, the capacity to care and feel emotions, as well as to have degrees of freedom and creative choice, are also relevant for moral standing. But these capacities all appear to be about grading moral standing, and not about grounding it.

Finally, intelligence is not subject to an intelligence-ism. Intelligence is the widest net. Outside of intelligent entities nothing morally matters. Within the realm of intelligent entities some matter more than others, but not just in virtue of being more intelligent. Other capacities grade moral standing. Schneider’s book explores the ubiquity of intelligence in the universe. We think it makes more sense to take it as the ground of moral standing.

San Jose State University
USA
anand.vaidya@sjsu.edu

O.P. Jindal Global University
India
rkrishnaswamy@jgu.edu.in

References

- Birch, J. 2020. The search for invertebrate consciousness. *Nous* 56: 133–53.
- Bradford, G. 2022. Consciousness and welfare subjectivity. *Nous* 57, doi: [10.1111/nous.12434](https://doi.org/10.1111/nous.12434).
- Chalmers, D. 2022. *Reality+: Virtual Worlds and the Problems of Philosophy*. New York: Norton & Company. 11.5
- Kammerer, F. 2022. Ethics without sentience. *Journal of Consciousness Studies* 29: 180–204.
- Schneider, S. 2019. *Artificial You: AI and the Future of Your Mind*. Princeton: Princeton University Press. 11.10
- Schneider, S. and E. Turner. 2017. Is Anyone Home? A Way to Find Out If an AI has Become Self-Aware. <<https://blogs.scientificamerican.com/observations/is-anyone-home-a-way-to-find-out-if-ai-has-become-self-aware/>>, last accessed 1 January 2024.
- Schwitzgebel, E. and H. Shevlin. 2023. Is it time to start considering personhood rights for chatbots? *Los Angeles Times*.
- Shepard, J. 2023. Non-human moral status: problems with phenomenal consciousness. *American Journal of Bioethics Neuroscience* 14: 148–157. 11.15
- Singer, P. 1975. *Animal Liberation*. New York: Harper Collins.
- Singer, P. 2009. *Animal Liberation*. New York: Harper Perennial Modern Classics.
- Udell, D.B. and E. Schwitzgebel. 2021. Susan Schneider’s proposed tests for AI consciousness: *promising but Flawed*. *Journal of Consciousness Studies* 28: 121–44. 11.20
- Wolfram, S. 2018. What is Intelligence? <<https://www.stephenwolfram.com/questions/2018/05/29/what-is-intelligence/>>, last accessed 1 January 2024. 11.25
- 11.30
- 11.35
- 11.40
- 11.45