

Retiring the “Cinderella view”: the spinal cord as an intrabodily cognitive extension

Marco Facchin [corresponding author]

Istituto Universitario di Studi Superiori IUSS Pavia
Department of Human and Life Sciences; Pavia, Italy
Palazzo del Broletto, Piazza della Vittoria n. 15, 27100 Pavia
marco.facchin@iusspavia.it
<https://orcid.org/0000-0001-5753-9873>

Marco Viola

Università degli studi di Torino
Department of Philosophy and Educational Sciences; Turin, Italy
Via Sant’Ottavio 20, 10124 Torino
marco.viola@unito.it
<https://orcid.org/0000-0001-7970-5406>

Elia Zanin

Rehabilitation Hospital at High Level of Specialization of Motta di Livenza, Treviso,
Italy
Spinal Cord Injury and Severe Acquired Brain Injury Unit
via Padre Bello 3/c, 31045 Motta di Livenza (TV)
elia.zanin@ospedalemotta.it
<https://orcid.org/0000-0003-1784-1933>

Abstract:

Within the field of neuroscience, it is assumed that the central nervous system is divided into two functionally distinct components: the brain, which does the cognizing, and the spinal cord, which is a conduit of information enabling the brain to do its job. We dub this the “Cinderella view” of the spinal cord. Here, we suggest it should be abandoned. Marshalling recent empirical findings, we claim that the spinal cord is best conceived as an intrabodily cognitive extension: a piece of biological circuitry that, together with the brain, constitutes our cognitive engine. To do so, after a brief introduction to the anatomy of the spinal cord, we briefly present a number of empirical studies highlighting the role played by the spinal cord in cognitive processing. Having so done, we claim that the spinal cord satisfies two popular and often endorsed criteria used to adjudicate cases of cognitive extension; namely the parity principle and the so-called “trust and glue” criteria. This, we argue, is sufficient to vindicate the role of the spinal cord as an intrabodily mental extension. We then steel our case considering a sizable number of prominent anti-extension arguments, showing that none of them poses a serious threat to our main claim. We then conclude the essay, spelling out a number of far-from trivial implications of our view.

Keywords: Spinal cord, Spinal cord injury, *sfMRI*, Extended mind, Embodied cognition, Internalism.

Link to the read only version: <https://rdcu.be/cxMmo>

Formatted in cosmic sans to be more easily readable from people suffering from dyslexia.

Declarations: Not Applicable.

Founding: MF has been funded by the PRIN Project "The Mark of Mental" (MOM), 2017P9E9N, active from 9.12.2019 to 28.12.2022, financed by the Italian Ministry of University and Research.

Conflict of Interests/competing interests: The authors declare no conflict of interests.

Availability of data and Material: Not applicable.

Authors' contribution: All authors contributed equally to the present manuscript.

Acknowledgements: We wish to thank the two anonymous reviewers for their insightful comments on an original version of the manuscript.

Retiring the “Cinderella view”: the spinal cord as an intrabodily cognitive extension

1 - Introduction: the Cinderella of the central nervous system

Embodied and extended approaches to cognition suggest that some constituents of the cognitive machinery lie outside the nervous system (Colombetti and Zavala 2019; Boem *et al.* 2021) or the agent's body (e.g. Clark 2008). Conversely, their critics endorse contingent intracranialism, claiming that, as a matter of contingent fact, the cognitive machinery is skull-bound (e.g. Adams and Aizawa 2008). Yet, anatomically speaking, extraneuralism is not the antithesis of contingent intracranialism. Numerous *neural* structures lie outside the cranium. Even ignoring the peripheral nervous system (but see Aranyosi 2013), the central nervous system consists of two components, one of which laying outside the cranium; namely, the spinal cord (SC). So, there is a third position between extraneuralism and (contingent) intracranialism: extracranialism. Extracranialism claims that the cognitive machinery is not exclusively located in the cranium, without *thereby* necessarily endorsing extraneuralism. In this way, it respects the neurocentric intuition many share.

Importantly, conceiving the SC as a cog in the cognitive machinery challenges the standard view according to which the SC is “The Cinderella of the Central Nervous System”, as Wolpaw and Tennissen (2001) poetically wrote. According to this “Cinderella view”, the SC is just a cable attaching the brain to the sensory periphery, which only realizes (cognitively irrelevant) reflexes (e.g. Kandel *et al.* 2012, pp. 790-810).

We think the “Cinderella view” ought to be abandoned. The cognitive consequences

of spinal cord injury (SCI), and the surprising data gathered through spinal functional magnetic resonance (sfMRI) strongly suggest that the SC is best conceived as an intrabodily cognitive extension: a piece of machinery that, together with the brain, conspires to produce a number of cognitive outputs and intelligent behaviors.¹ That is what we will claim in the following.

Two caveats before proceeding. First, we will identify the cognitive *in extension*, by referencing paradigmatically cognitive processes such as memory, learning, integrating information, and so forth. We will not endorse any specific mark of the cognitive (i.e. a set of individually necessary and/or jointly sufficient criteria identifying all and only cognitive processes). We acknowledge that the lack of a mark of the cognitive might "muddy" our decision on whether to call some process cognitive in contested cases (e.g. do bacteria cognize? see Adams 2018). But the processes we will discuss are all paradigmatically and uncontestedly cognitive, so we are safe from "muddying".

Secondly, and relatedly, since the evidence we will present has been communicated using a representationalist lexicon, we will speak in representationalist terms. This does not mean, however, that we are committed representationalists, nor that we think that cognition must *necessarily* involve representations (see Ramsey 2017). Our usage of the representationalist lexicon, thus, might be seen as a purely pragmatic matter.

We structure our paper as follows. In the next section, we provide a succinct introduction to spinal anatomy, and review some recent threads of evidence

¹ Allen *et al.* (2009) made some steps along this same path providing an empirical informed philosophical analysis of the rat's SC and of its role in some cognitive processes. Yet they do not claim that the SC is a constituent part of the cognitive machinery.

suggesting that the SC plays an active role in cognitive processing. In section 3, we will articulate our philosophical claim, presenting a *parity argument* for the SC as an intrabodily cognitive extension. In section 4, we will confront a number of counterarguments that can be levelled to our claim. Section 5 concludes the essay by pointing out some far from trivial implications of our claim.

2 - Spinal contributions to cognitive processing: some evidence

The SC is an anatomically defined structure of the central nervous system, resting in the spinal canal of the vertebral column. It extends from the *medulla oblongata* in the brainstem to the *cauda equina* in the lumbar vertebrae. The SC is anatomically divided into five regions (cervical, thoracic, lumbar, sacral and coccygeal), each organized into segments, which are anatomically defined by the roots of the departing fascicles (see figure 1).

[Figure 1 Omitted. Please, refer to the published version of this article.]

The SC reverses the anatomical arrangement of the brain, insofar cell bodies (grey matter) are surrounded by myelinated nerve fibers (white matter). The grey matter core is divided into two horns. The dorsal horn is the ascending pathway, carrying information from the transducers at the sensory periphery (such as tactile receptors, nociceptors and thermoreceptors) to the brain. Conversely, the ventral horn is the descending (brain-to-periphery) pathway, and it is typically considered the locus of production of both voluntary movements and automatic reflexes, as it contains the motor neurons that directly innervate the muscles (Kandel *et al.* 2012, Ch. 15; 35).

Lesions of the SC are called spinal cord injuries (SCIs). SCIs might have traumatic

or non-traumatic origin, but all bring about changes to the sensory, motor and autonomic function of the SC, which might be temporary or permanent, depending on the severity and level of lesion (Holtz and Levi 2010). Two common outcomes of SCI are paraplegia and tetraplegia. In both cases, depending on the level of lesions, the patient loses some sensory and motor functions and might suffer from autonomic dysfunctions (i.e. bladder, bowel, sexual dysfunctions). In paraplegia, the loss of sensorimotor function is confined to the lower limbs, as a result of thoracic, lumbar or sacral lesions. Tetraplegia also affects the upper limbs, and it is usually the outcome of cervical lesions (Kandel *et al.* 2012, Ch. 34-35; Holtz and Levi 2010).

Thus presented, the SC appears to be just a cable connecting the brain and sensory periphery, allowing the former to control the latter, just as the "Cinderella view" contends. Yet, several recent threads of evidence put this view under pressure. Paralleling the investigation of cerebral cognitive functioning, the study of SC functioning has been historically accomplished by correlating macro-structural lesions with resulting sensori-motor impairments. Imaging data from the SC have traditionally been difficult to acquire because of many thorny technical and physiological issues, such as the non-uniform magnetic field of bones and tissues near the SC, the physiological motion of the SC, and its small dimensions. Yet, technical advances in functional neuroimaging techniques have recently allowed researchers to overcome these pitfalls (see Stroman *et al.* 2014). These advancements, which generated *spinal cord* functional magnetic resonance imaging (sfMRI) techniques, allowed researchers to explore SC functioning *in vivo* (Wheeler-Kingshott *et al.* 2014). Here, we present some recent empirical evidence suggesting that the SC is a part of the cognitive circuitry.

Let us start in the proximity of the sensory and motor functions traditionally ascribed to the SC. Thus, consider an agent's *peripersonal space*: the representation of the space immediately surrounding an agent's body. The representation is centered upon the agent's effectors (Rizzolatti 1997; Di Pellegrino and Làdavas 2015), and its realization involves a fronto-parietal network tasked with integrating different sensory modalities (Brozzoli *et al.* 2011). Such an integration allows objects falling into an agent's peripersonal space to be represented as affordances; that is, in terms of the actions they invite (e.g. Costantini *et al.* 2011). Importantly, although the representation of the peripersonal space is related to motor control, it is not a motor command. Typically, motor commands specify actions in intrinsic (muscle-based) terms, such as joint torques and muscle forces (e.g. Hollerbach 1982; Todorov 2004), none of which seems represented in the peripersonal space. Thus, the representation of the peripersonal space is "sandwiched" between input reception and action execution (Hurley 2001), and qualifies as a *bona fide* cognitive process even according to classical cognitivist standards.

The "Cinderella view" would thus predict that the SC plays no role in the representation of the peripersonal space. Recent empirical data, however, suggest otherwise. The Crossmodal Congruency Task (CCT) described in (Scandola *et al.* 2016; replicated in Scandola *et al.* 2020) provides a nice example. In CCTs, subjects have to judge the height (high vs low) of a tactile stimulus delivered to their hands while ignoring a visual distractor. When stimulus and distractors are applied ipsilaterally, the subjects' reaction times increase, but only if the subjects represent both within their peripersonal space. This is because unlike the extrapersonal (far) space, the peripersonal space is multimodally represented, allowing visual distractors to

interfere with tactile processing. Scandola and colleagues harnessed this well-known effect to assess how paraplegic patients² represent the space near their disconnected limbs. Strikingly, whereas the subjects' reaction times increased when both stimulus and distractor were ipsilaterally presented near their hands, the increase of reaction times vanished when the distractor was placed near the subjects' feet, suggesting their peripersonal space has *contracted selectively*, excluding the disconnected limbs.³ It thus seems that physical damage to the SC affects the representation of the peripersonal space.

This *selective* effect of SCI might perhaps be accounted for by the cortical changes triggered by the prolonged immobility of (and concomitant lack of sensory feedback from) the affected limbs (see Scandola *et al.* 2016). But SCIs also have *non selective* effects on the peripersonal space. Hence, consider reachability judgments: in certain experimental tasks, subjects are asked to decide whether a target object falls within the subject's reach. A subject's answers to these questions provides a popular indirect measure of a subject's peripersonal space, as a subject's peripersonal space ends where objects stop to be graspable (e.g. Costantini *et al.* 2010)

Importantly, these judgments are subtly, but systematically, altered in paraplegic patients (Sedda *et al.* 2019). To begin with, the reachability judgments of healthy subjects exhibit a systematic bias, as the reachability range is typically over-estimated (see Ambrosini *et al.* 2012). Injured subjects, however, did not manifest any such bias. Moreover, when formulating reachability judgments, the reaction times of healthy subjects drops off as a function of object proximity: the

² In all clinical data we discuss, SCI participants *did not* suffer any concomitant brain damage, nor suffer from any psychological disorder.

³ Importantly, the data presented in (Scandola *et al.* 2020) also suggest that the lower the lesion, the more the peripersonal space selectively contracts around feet.

closer (or further away) the object, the quicker the judgment. But no such effect was found in SCI patients. Lastly, whereas healthy subjects sharply transition from positive to negative judgments when the objects are presented beyond a critical threshold, paraplegic patients exhibit a more graded transition, suggesting their peripersonal space has less defined boundaries. Interestingly, these differences were more marked in SCI patients with more severe lesions.

Notice all these effects are not selective: reaching is a hand action; and paraplegia does not affect the upper limbs. As a consequence, these findings suggest that physical damage to the SC affects the representation of the peripersonal space in a way that cannot be "explained away" just by invoking the lack of feedback from the affected limbs.

Perhaps the representation of the peripersonal space is *too peripheral* to deserve to be called a "real" cognitive process. It might be too close to action to put any significant pressure on the "Cinderella view". Imagination, however, surely is not so peripheral. Hence, discovering that the SC is involved in imagination would put the "Cinderella view" under pressure. Recent empirical evidence speaks of such a role, at least when it comes to motor imagery (MI; see Di Rienzo *et al.* 2014 for a review).

MI involves the usage of one's motor skills *offline*, to imagine executing a given motor action (Jeannerod 1994). Importantly, since MI recruits the same regions subserving the execution of *online* actions (e.g. Alkadhi *et al.* 2005; Miller *et al.* 2010; Olsson 2012) it is constrained by the same biomechanical constraints of real actions. This clearly emerges in laterality tasks, in which subjects must judge the laterality (left or right) of a rotated body part by mentally rotating their corresponding effector. The harder the target position to reach, the slower and more prone to

error subjects become (Parsons 1987; 1994). No such effect has been detected in cases of purely visual imagery, which in fact involves a different set of cortical areas (e.g. Albers *et al.* 2013).

No such effect is also detected when paraplegic⁴ and tetraplegic subjects engage in laterality tasks, forcing some researchers to conclude that SCI patients are unable to resort to MI to complete laterality tasks (Fiori *et al.* 2014). As strange as this might sound, the empirical finding seems to be robust. In fact, it is supported by studies on patients suffering from locked-in syndrome (Conson *et al.* 2008; 2010). Locked-in syndrome is a disruptive condition in which damage to the brainstem prevents the usual message-passing from brain to SC (and *vice versa*), resulting in the patient being deafferented from her whole body, with the partial exception of the eyelid. Strikingly, the results obtained from patients suffering from locked-in syndrome engaged in laterality tasks do not reflect the effects of biomechanical constraints too, as if the contribution of the SC were *needed* to engage in MI.

These findings on the role of the SC in MI seem to suggest that the SC plays a role in cognitive processing. It is worth noting, at this juncture, that the SC is not a "static" structure, hosting only a handful of hard-wired reflexes. To see why this is the case, consider the dual role of spinal plasticity. First, it allows the SC to learn, and directly execute when appropriate, *motor reflexes*. This allows large chunks of behavioural control to be automatized, lowering the computational burden motor control places onto the motor cortex (Wolpaw 2007; Lungu *et al.* 2010). Secondly, spinal plasticity also enables, and actively contributes to, the learning of complex motor sequences. Hence, the SC appears to contribute to our procedural memory. As

⁴ In the case of paraplegia, the effect might be limited to the deafferented limbs, see (Ionta *et al.* 2016).

a nice example, consider the role the SC plays in *implicit motor learning*: the form of practice-based learning enabling us to seamlessly execute complex motor sequences such as playing Bach on the piano or to fluently typing on a keyboard.

Implicit motor learning is often evaluated through *serial reaction times* tasks (Robertson 2007), where participants must rigidly respond to a closed set of stimuli while their reaction times are being measured. As the order of stimulus presentations is predictable⁵, the decrease of reaction times across repeated trials indicates that the subjects are learning the motor sequence needed to efficiently respond to the stimuli. No such decrease in reaction times, however, was detected when paraplegic patients were engaged in a serial reaction time task (Bloch *et al.* 2016), suggesting a deficit in implicit motor learning. Importantly, participants did not even show an *increase* of reaction times when a non-predictable series of stimuli were intermixed in the trials. Such an increase is expected (and was detected in the control group of uninjured subjects) due to the interference effect between the implicitly learned (predictable) sequence and the non-predictable sequence. It thus seems that SCI prevents the (implicit) learning of the sequence, rather than just a decrease of reaction times.

These findings are nicely complemented by a *sfMRI* study by Vahdat and colleagues (2015), which showed that healthy subjects exhibit significant clusters of spinal activity during implicit motor learning tasks. In this study, subjects were scanned in two conditions. In the first, they had to press the buttons of a pad so as to act out a very simple motor sequence, which frequently recurs in normal (ecological) contexts and that they thus did not have to learn. In the second

⁵ Which does not entail that participants notice its predictability.

condition, they had to press the buttons acting out a fairly complex motor sequence that they most likely had to learn during the task. The imaging data showed the presence of two spinal clusters of activity in the predicted cervical level. These clusters of activity, however, had different coordinates in the two conditions, and the amplitude of the BOLD⁶ signal was much larger in the second condition. Moreover, albeit the reaction times of the subjects decreased in both conditions, only in the second condition the decrease of reaction times correlated with changes in the spinal BOLD signal. More strikingly still, Vahdat and colleagues found that these clusters of activity accounted for 24% of the total variability of the BOLD signal, and that such activity was mostly *independent* from concomitant brain (cortical and subcortical) activity. In fact, 81% of the spinal activity was neither positively nor negatively correlated with brain activity. In summary, these data strongly suggest that, in implicit motor learning, spinal activity is *not* a mere reverberation of cortical (or subcortical) activations, and that such an activity plays a large and important role in implicit motor learning.

Importantly, the SC is not exclusively involved in action-related cognitive processing. It also appears to be involved in perceiving and expressing affective states.

Several studies suggest that SCI correlates with a reduction in affective processing. An early correlation was found by Montoya and Schandry (1994), who reported that, compared to healthy controls, SCI patients report overall less anxiety and less emotional experience. Pistoia and colleagues (2015) provide similar data,

⁶ Most fMRI studies are based on the Blood Oxygenation Level Dependent (BOLD) signal: a signal sensitive to the changes of magnetic properties of blood due to neurovascular activity, which offers an indirect measure of neural activations in small neural regions (Logothetis 2002).

reporting that, compared to healthy controls, SCI individuals showed difficulties to recognize fear- and anger- inducing pictures, in a way that directly correlates with the height of the lesion. In a similar experiment, Guadagni and colleagues (2019) found results suggesting that SCI patients tend to rely more on *cognitive* empathy rather than emotional simulation to judge the valence of an emotional stimulus.

Complementarily, the development of *sfMRI* led to the discovery of important clusters of *SC* activity during affective processing. Thus, Smith and Kornelsen (2011; Kornelsen *et al.* 2015) found different clusters of cervical and thoracic spinal activity when healthy subjects were exposed to emotionally valenced stimuli, especially negative valenced ones. Different patterns of *SC* activity were also found when subjects listened to aversive (compared to neutral) sounds (Smith *et al.* 2018a). Smith and colleagues (2018b) also investigated the *SC* role in the *production* of emotional expression. Their findings indicate that the thoracic level of the *SC* is significantly involved in the production of disgusted facial expressions (contrasted with neutral ones), corroborating the hypothesis that disgusted facial expressions might be part of a motor program aimed at reducing air intake.

These recent findings are suggestive of a role for the *SC* in affective processing, although the nature and the extent of this contribution has still to be empirically determined.

3 - The spinal cord as an intrabodily extension: the parity argument.

The previous section exhibited several distinct threads of evidence motivating the rejection of the "Cinderella view". Far from being a mere cable keeping the brain attached to the body, the *SC* appears to play a significant role in a number of

cognitive processes. We suggest that the evidence provided above forces us to consider the SC as an *intrabodily cognitive extension*: a piece of neural machinery that, together with the brain, constitutes a subject's cognitive system.

To do so, we show that the SC satisfies the two criteria by means of which claims of cognitive extension should be adjudicated, namely the *parity principle* and the *trust and glue* criteria (Clark and Chalmers 1998; Clark 2008; 2010).⁷

Consider, first, the parity principle:

Parity Principle: if, as we confront some task, a part of the world function as a process which, were it to go on in the head, we would have no hesitation in accepting as a part of the cognitive process, then that part of the world is (for that time) part of the cognitive process (Clark 2008: 77).

The principle is typically applied as follows (e.g. Clark and Chalmers 1998: 13). First, one considers a *candidate* distributed or extended problem solving system (such as Otto and his notebook), and identifies the functional contributions made by external elements (e.g. the role of the notebook in storing information). Then, one imagines a purely inner system solving the same task; that is, a *cerebral* system or sub-system in which the functional contributions previously made by the external elements are now made by appropriate cerebral components (such as Inga's biological memory).⁸ If these internal elements would count as genuinely constituent parts of the cognitive machinery in the internal case, then we *should* count their external counterparts as constituent parts of the cognitive machinery too, and the *candidate* distributed problem solving system as a *real* extended problem solving system.

⁷ As repeatedly pointed out in the literature on the extended mind (e.g. Gallagher 2018; Kirchhoff and Kiverstein 2019), not *all* claims of cognitive extension need to be adjudicated by means of these two criteria. But these two criteria are the only ones immediately relevant for our claim here: “second wave” criteria based on manipulation and “third wave” criteria based on cultural scaffolding do not seem to apply to the case of cognitive extension we are here interested in.

⁸ Notice that the functional parity of inner and outer resources is *inbuilt in the scenario the parity principle asks to imagine*. If internal and external resources are not imagined to be functionally on a par “one doesn't have a failure of parity; one simply hasn't set things up properly” (Wheeler 2019: 86).

This is supposed to be a "rule of thumb" to determine the constituents of the cognitive machinery in a way that ignores "old metabolic boundaries of skin *and skull*" (Clark 2013: 195, emphasis added). It expresses a straightforward insight: what matters, in order for something to qualify as (part of the) material basis of a cognitive process is neither its spatial location relative to skin and skull nor its gross physical substratum. What matters is instead its *functional and computational poise*: the way in which it supports an agent's problem solving routines leading the agent to behave intelligently in regard to the task at hand (see Clark 2008: 76-82).

Let us apply the parity principle to the case at hand. In section 2, we have briefly described the role the SC seems to play in various cognitive processes. Hence, we have a candidate extended cognitive system, composed by the SC and the brain. Now consider, as the parity principle requires, the following counterfactual scenario in which the contributions previously made by the SC are made by a cerebral component (Wheeler 2019). Imagine that a bit of the cortex (let us dub it *yet another motor cortex, YAM*) was largely supposed to be a cortical module for the execution of reflexes. Suppose now that sound empirical evidence suggests the following: that YAM increases its neurovascular activity when sequences movements are learned; that its plasticity is found in part responsible for the learning of these sequences, and that focal lesion to the YAM impairs motor learning, motor imagery, and substantially modify affective processing and how an agent represents its peripersonal space.

Neuroscientists would likely conclude that YAM is no "module for reflexes" at all, but that it instead is *part of one or more larger networks* performing a variety of cognitive processes. After all, neuroscience is growingly placing emphasis on the

distributed nature of most cognitive processes (Bressler and Menon 2010). Lesion studies and early day neuroimaging often proposed simple mappings between cognitive processes and neural regions, e.g. by locating fear into the amygdala; but these simple localizations are best construed as heuristics (McCauley and Bechtel 2001). Starting from here, neuroscientists pursue their investigation so as to uncover the network subserving a specific process. In a way, this suggest an intraneural cognitive extension: for instance, fear processing has been extended beyond amygdala so as to include the pulvinar and some fronto-occipital cortices, plus other cortices depending on the explicit/implicit nature of processing (Tao *et al.* 2021). Notice that such an extension pertains to the ordinary practice of present-day neuroscience. Now, take disgust. It used to be localized within the anterior insula. Suppose that YAM consistently activates when subjects produce disgust expressions. There seems to be no principled reason not to extend the disgust network so that it includes YAM, just like the pulvinar was included in the fear network. But, YAM is just the fictitious counterpart of another neural, albeit non-cortical, region; namely, the SC. Indeed, evidence of activation of SC in disgust expressions (Smith *et al.* 2018b), just like in all the cognitive processes assigned to YAM, is documented at length in the previous section.

So, if YAM would be considered *part of* one or more larger networks subserving cognitive processing, why the SC shouldn't? They are, by construction, functionally identical. The answer cannot just be "because the SC is not cortex": neither are the amygdala or the pulvinar, and yet nobody denies their status in implementing cognitive processes. Hence, if the parity principle holds, and our story about YAM is plausible (as it seems), then we *should* conclude that the SC is *part of* one or more networks

that perform cognitive processes, in the standard sense in which the amygdala or the pulvinar are parts of similar networks. Perhaps it might be objected that the parity principle alone is *too lax* a criterion to adjudicate claims of cognitive extension. In fact, it should be complemented by the "trust and glue" criteria (Clark 2010). According to these, a putative cognitive extension identified through the parity principle should also be (i) readily available and typically invoked in actual cognitive processing, (ii) store information automatically endorsed upon retrieval, and (iii) store information easily accessible when required.

We believe that the SC easily satisfies (i) to (iii) in conjunction. To begin with, as illustrated above, the SC is involved in a number of different processes, so it seems correct to conclude that it is typically invoked in cognitive processing. It also seems readily available: unless the brainstem is seriously lesioned, nothing seems to prevent the SC from playing the required role when needed. It thus seems correct to conclude that the SC satisfies (i)

It seems also correct to say it satisfies (ii) and (iii). After all, as seen above, the SC at least partially stores our *procedural knowledge*: the kind of know-how that is requested to skillfully complete motor tasks. And it seems correct to say that we rely on such knowledge *unreflectively*: we simply deploy our relevant know-how when needed, without assessing the epistemic quality of the information constituting it. Hence it seems correct to say that the relevant information stored in the SC is automatically endorsed upon retrieval, as (ii) requires. Moreover, our procedure knowledge is *just there* as and when needed, as (iii) requires. Here, our focus has been on procedural knowledge just to make the point we are trying to make vividly appear. Analogous considerations hold for MI or the representation of our

peripersonal space: we just engage in acts of MI when needed, and we do so *without* monitoring what our SC (or, equivalently, our motor cortices) are doing. In a similar spirit, we *just see* that an object is within our reaching range, or alter the air intake of our lungs when we are disgusted, without scrutinizing what sort of processes are happening inside us. It thus seems the SC satisfies (i)-(iii) in the exact same way the encephalon does. On the fairly uncontested presupposition that the brain satisfies (i)-(iii) and it is thus part of our cognitive system, we conclude that the SC is also part of it.

The parity argument we just made might seem far too glib. A sceptical reader might, for instance, point at the impressive number of arguments challenging claims of cognitive extension (see Menary 2010; Clark 2013, pp. 192-212) and point out that our claim cannot be *that easy* to make and defend. Agreed. In the next section, we will consider, and defuse, the most prominent objections that could be levelled at our claim.

4 - Defending the parity argument: some objections considered

The "extended mind thesis" has been the target of a variety of counterarguments. Here, we will consider, and defuse, the ones to which our claim seems susceptible to the most. To keep things in good order, we encapsulate the discussion of each counterargument in a dedicated subsection.

Two *caveats* before moving forward. First, we will *not* commit here to any criterion to determine the "boundaries of the mind". We just aim at showing that, of the existing ones, none provides reasons to leave the SC out.

Secondly, we will not say that the SC "is cognitive". In the past, the usage of the

"is cognitive" locution when dealing with claims of cognitive extension give rise to misunderstandings, as if claims of cognitive extensions were claims about extra-cerebral objects being thinkers or cognizers *on their own*⁹ (see Adams & Aizawa 2010: 67), rather than claims about objects being constituent parts of cognitive systems (see Clark 2008: 86-93). Since we wish to avoid this kind of misunderstanding, we will avoid saying that the SC "is cognitive". We will, however, defend the claim that (barring pathological conditions involving severe SCI) the SC is a constituent part of the human cognitive system.

4.1 - The cognitive bloat objection

The cognitive bloat objection is basically a slippery slope objection against claims of cognitive extension (Rowlands 2010: 93-95). It starts by granting, for the sake of argument, some extension claim, such as that entries in a notebook might constitute a subject's beliefs (e.g. Clark and Chalmers 1998). It then notices that, *if* in the accepted case the cognitive extension claim is true, *then* it must also be true in other cases. For instance, if entries in a notebook qualify as "extended beliefs", then entries in a phone book (Rupert 2004) or encyclopedia (Sprevak 2009) qualify too. But then we are forced to ascribe subjects an inordinate amount of beliefs, which the subject does not plausibly hold (Ludwig 2015). Hence, either claims of cognitive extension lead to absurd consequences, and in this case claims of cognitive extension are denied by *reductio*, or are simply false (if the consequent is false, the antecedent must be false too).

We think that our claim is not threatened by the cognitive bloat objection for at

⁹ We suspect that this partially prevented (Allen et al. 2009) from making the claim we defend here.

least two reasons.

First, in order for the "cognitive bloat" to go through, one must be committed to a claim about cognitive extension that allows for *progressively weirder* items to be counted as constituents of a subject's cognitive circuitry. But the claim of extension we are defending does not seem to allow that. The jump from "the SC is part of a subject's cognitive apparatus" to "an encyclopedia is part of a subject's cognitive apparatus" seems a big, and poorly justified, leap. This is because the claim of extension we are defending affords us a number of *principled criteria* which can block the rampant expansion of the mind into the world. For instance, we are free to adopt a neuro-chauvinistic position, according to which only neural stuff can possibly constitute the mind. That would block the bloat. Surely, doing so is inconsistent with the "spirit" of cognitive extension, but, crucially, *is not inconsistent with what we have argued thus far*. Or we could adopt a commonsensical position according to which only what is "inside" of the boundaries drawn by perception and action counts as a constituent of a subject's cognitive system (see Chalmers 2008; 2019). Again, that would block the bloat, while allowing the SC to count as a constituent. To restate our first caveat, we here do not wish to endorse any of these two options: we are just pointing out that the claim we are defending allows us to effectively block the bloat objection.

Secondly, the bloat objection is a menace only insofar it forces one to ascribe new (and exotic) *propositional attitudes* to a subject (see Rupert 2004; Ludwig 2015); the *reductio* the bloat objection brings about hinges upon these exotic ascriptions. But accepting that the SC is a constituent of a subject's cognitive machinery does not seem to license the ascription of such exotic propositional attitudes. The data

examined in section 2 do not seem to have any entailment in regard to what a subject believes (or hopes, desires, etc.). It thus seems that the claim of cognitive extension we are defending is safe from the *reductio* the bloat objection threatens to bring about.

4.2 - The cognitive systems objection

The cognitive bloat objection highlights the need to identify the boundaries of our cognitive machinery. The cognitive system objection identifies them as follows: a state, process or mechanism: "is cognitive if, and only if, it is a state of, or process occurring in, mechanisms that are elements of the integrated set members of which contribute causally and distinctively to the production of cognitive phenomena." (Rupert 2009: 35; see also Rupert 2004; 2010). Since extraneural extensions do not belong to any such integrated system, they are not part of our minds - or so Rupert suggests.

Prima facie, the objection appears toothless against our claim. For the SC and the brain jointly form an anatomically defined system, the CNS. According to Rupert's definition, however, cognitive systems need not be identical to anatomically defined systems. Consider neuroglial cells: they surely are part of the central (and peripheral) nervous systems, but they are not typically held to be responsible for the production of cognitive phenomena. Hence, they are not part of any cognitive system, according to Rupert's proposed definition.¹⁰ However, the same is hardly true of the SC. In fact, the evidence exposed in section two suggests that the SC significantly

¹⁰ They surely do not belong to the cognitive system according to Adams and Aizawa (2010: 593, fn. 23): "It is, of course, an oversimplification to say that the whole of the brain realizes cognitive processes, since there may well be glial cells or blood vessels or other such structures that do not". However, glial cells' involvement in cognitive functions is currently being debated (see Butt and Verkhatsky 2018).

contributes, together with cerebral mechanisms, to the production of different cognitive outputs; just as Rupert (2009: 37-44) requires.

This verdict might be challenged by the *screen-off* criteria Rupert (2009: 43) proposes to correctly identify cognitive systems. A first criterion excludes mere sources of input to the cognitive system. Thus, if a putative cognitive mechanism seems to frequently conspire with others to produce cognitive phenomenon *only because* it causally affects a limited subset of them, it should not be counted as part of the cognitive system, but merely as an environmental source of input the system relies on (Rupert 2009: 43). Consider, for instance, *glasses*. Short sighted people often rely on them, and they do causally contribute to the production of cognitive outcomes (e.g. visually guided action). Yet they are not, intuitively speaking, part of our cognitive system. Rupert's screen-off criterion allows us to secure this intuition, letting us regard our glasses as a mere source of input to the real cognitive system. A similar criterion could be invoked to screen-off the *behavioural products* of cognitive processing. That is, if a mechanism seems highly integrated with other cognitive mechanisms *only because it is affected* by a limited subset of these, it should not be considered part of the cognitive mechanism, but a mere behavioural correlate of the system's functioning. Consider, for instance, *lip movements*. Surely lip movements causally contribute to all human oral verbal expression. Yet they are not, intuitively speaking, part of our cognitive machinery. Applying this second screen-off criterion would secure our intuition (we owe this example to Krickel 2019).

Do these screen-off criteria exclude the SC? The corticospinal tract projects to multiple brain areas, being connected not only with primary somatosensory and primary motor cortices, but also with premotor and supplementary motor areas, as

well as somatosensory areas in the parietal lobe (Guyton and Hall 2005: 687; Kandell *et al.* 2012: 843). Thus the SC appears well integrated with the rest of the cognitive system, as the screen-off criteria require.

Furthermore, there are cases in which the SC activity cannot be *just* the brain's conduit of input and output. Consider again the evidence regarding MI exhibited in section 2. Recall: subjects were asked to judge the laterality (left or right) of an effector, by mentally simulating the rotation of the corresponding effector to the target position. This is a well defined cognitive task, with well specified stimuli and responses. The stimuli are rotated images of bodily parts. Hence the input the brain receives is *visual*. Thus it reaches the brain travelling from the optic nerve to the Lateral Geniculate Nucleus, and not the SC. Consider now the response: it typically is verbal. So it requires mouth movements, controlled through the facial nerve.¹¹ It thus seems in these tasks the SC cannot be *just* providing input to the brain and *just* executing the brain commands, for the input to and the output from the brain *need not travel there*.

So, there is *at least one* cognitive task in which the SC seems to play some important computational role, which cannot be *just* the role of being a conduit for inputs and outputs. Accordingly, we conclude that Rupert's screen off criteria should not screen the SC off.

4.3 - The mark of the cognitive

Given our purpose here, the most straightforward way to introduce the idea of a

¹¹ In three laterality tasks we cited (Conson *et al.* 2008; 2010; Fiori *et al.* 2014) subjects had to respond through eye movements. This, however, does not put pressure on our point, as eye movements are controlled through cranial nerves too.

mark of the cognitive is in terms of a set of individually necessary and jointly sufficient conditions identifying all and only cognitive phenomena. No such mark has yet been proposed. However, some have suggested some necessary (Adams and Aizawa 2001; 2008; 2010) or jointly sufficient (e.g. Rowlands 2010), conditions.¹² These proposals put forth candidate parts of a mark of the cognitive, which enable us to adjudicate (at least *pro tempore*) whether a given phenomenon (i.e. item or process) qualifies as cognitive. It is thus immediately clear how a (part of the) mark of the cognitive hinges on the present issue: by showing that the SC satisfies the relevant (parts of the) mark, one shows that the SC is (at least) a candidate constituent of our *cognitive machinery*.¹³ Since in this section we want to defend our claim from possible objections, we consider Adams and Aizawa's (2001; 2008) proposed mark, as it is the one tailored to secure contingent intracranialism.

Adams and Aizawa propose two individually necessary conditions as parts of the mark of the cognitive. In their view, the cognitive is (a) individuated in terms of specific kinds of mechanisms and (b) involves non-derived content. We unpack each point in turn, showing that the SC satisfies it.

Condition (a) says that the cognitive is individuated by specific kinds of mechanisms. The point can be straightforwardly expressed as follows: there are a number of ways in which one can carry out a cognitive task. One, for instance, might pass the Turing test either by understanding the examiner's questions and responding appropriately, or by "fishing" the right answer from a large database containing all

¹² Notice, however, that (as a reviewer aptly noticed), nothing obliges the mark of the cognitive to be developed as a set of individually necessary and jointly sufficient conditions. One might also opt for an approach based on family resemblance, singling out paradigmatic cognitive phenomena and processes to then judge the extent to which extended cognitive phenomena and processes are similar to paradigmatic ones (e.g. Wheeler 2011).

¹³ Thus, albeit the mark of the cognitive is typically leveraged as an objection to the extended mind (e.g. Adams and Aizawa 2001), it can be fruitfully used as an argument supporting it (e.g. Rowlands 2010; Wheeler 2019).

possible human conversations (Adams and Aizawa 2001: 52). Only the first way to pass the Turing test counts as cognitive, and it does so *in virtue of* the specific mechanism that solves the cognitive task. So only some specific mechanisms can in principle count as cognitive. Which mechanisms? Apparently, the ones identified by our best cognitive psychology (see Adams and Aizawa 2001: 52; 2008: Ch. 4).

But if this is the case, then the processing occurring in the SC doubtlessly satisfies condition (a), for the kind of processing taking place in the SC is one of the objects of study of cognitive neuroscience/neuropsychology, which is a branch of cognitive psychology. The empirical evidence provided in section 2, we think, make this point perspicuous enough. Thus, if the cognitive is partially identified in terms of the kind of processes discovered by cognitive psychology, then surely the processes taking place in the SC will count as cognitive - provided it also satisfies condition (b).

Now, condition (b) says that cognition involves non-derived content; that is, states or items which bear content independently of some other already contentful state (Adams and Aizawa 2008: 9). Non-derived content is thus determined by non-contentful factors or relations, such as the ones typically involved by naturalistic accounts of content (Adams and Aizawa 2008: 31). Do the states of the SC carry such content?

To start, the SC is typically viewed as a representational system in contemporary cognitive neuroscience. Theories of motor control, for instance, view the SC as a representational system, busy transmitting *messages* from the motor system to the actuators. In fact, different theories of motor control hold that the SC transmits messages with different *contents* to the motor plant. For instance, according to optimal control theory (e.g. Todorov 2004; Bays and Wolpert 2007), the SC transmits

motor commands; that is, representations that specify the movements to be executed. Conversely, active inference says that the SC transmits sensory predictions: that is, representations of the sensory outcomes of movement (Friston 2011; Adams *et al.* 2013).

One might contend that these contents are merely *ascribed* to the SC by some external observer (say, a neuroscientist) for some purpose (say, simplifying the explanation of motor control). If that were the case, then the content of the representations in the SC would be derived, as it would depend on the already contentful interpretation of an observer. Hence, point (b) would not obtain.

It is relatively easy, however, to show that the states of the SC satisfy at least some naturalized accounts of content, thereby qualifying as representations with non-derived content. Thus consider so-called *teleo-informational accounts* of content (e.g. Dretske 1988; Neander 2017; Shea 2018, Ch. 4).¹⁴ Very broadly, according to these accounts:

A state R of a system S is a representation of a target T *iff*:

- (i) R indicates T; &
- (ii) R has the function of indicating T within S.

Since teleo-informational accounts of content are widely known, little unpacking seems needed. Indication is a natural relation of covariance between R and T, such that the tokening of R raises the probability of T being the case. The relevant notion of function mentioned by (ii) is that of teleological function; namely, the output an item is supposed to bring about in virtue of its history of (re)production. Thus, just

¹⁴ Adams and Aizawa are never explicit about what sort of naturalized account of content they favor. For this reason, we here consider what strikes us as the most popular one. Notice further that, due to space limitations, we will not consider each individual teleo-semantic account of content (in spite of their individual differences). Rather, we will focus on the rough idea underlying them all.

as hearts are *supposed* to pump blood given the evolutionary history of animals, Rs must be *supposed* to indicate Ts given the reproductive history of Ss.

To see how states of the SC satisfy (i) and (ii), consider the role the SC plays in pain perception (Kandel, *et al.* 2012: 530-536). The neurons of the dorsal horn of the SC participating in pain perception are highly segregated based on the inputs they receive. Lamina I neurons, for instance, receive inputs from nociceptors on the skin, whereas lamina V neurons receive inputs from nociceptors in the viscera. Moreover, these neurons are segregated within each lamina based on the *kind* of nociceptor from which they receive their inputs, and different kinds of nociceptors are selectively sensitive to different noxious stimuli (e.g. thermal nociceptors respond to extreme temperatures, whereas mechanical nociceptors respond to vigorous pressures). Thus, for instance, there is a special sub population of lamina I neurons responding selectively to intense cold temperatures.

It seems to us that these neurons satisfy (i) and (ii) in a fairly straightforward way. Their states (that is, their activation) robustly and selectively covary with a well-specified environmental parameter (e.g. intense cold), in a way such that observing the activation of these neurons *would* make us more certain that a cold stimulus has been applied to the portion of a subject's body which is monitored by the receptive field of these neurons. Hence, these neurons indicate the presence of noxious stimuli in the same unproblematic sense in which the selective activation of neurons in the primary visual cortex indicates the presence of specific kinds of stimuli (e.g. Hubel and Wiesel 1962; 1968). And it is extremely likely that these neurons have been selected by natural selection to perform a certain function, which can be supposed to be that of indicating the presence of a specific kind of external

noxious stimulus, such as intense cold. Indeed, such a supposition appears to be the *default* supposition when examining the nervous system (e.g. Kandel, *et al.* 2012, Ch. 21; Neander 2017, Ch. 4). If this is correct, then it seems that at least some states of the SC qualify as representations with non-derived content according to a prominent theory of content.¹⁵

At this juncture, Adams and Aizawa might further contend that albeit the SC merely *transmits* its contents, without manipulating them. Since cognitive processing involves the manipulation of content, the spinal cord is not involved in any cognitive process (Adams and Aizawa 2008: 17 make this point about the *corpus callosum*).

But that objection strikes us as flawed for two reasons. First, it assumes that the SC "merely carries" contents, without manipulating them. But this flies in the face of the empirical evidence we presented above. It also flies in the face of *at least some* computational accounts of neuronal functioning, according to which the SC is a computational engine (e.g. according to active inference, the way in which spinal alpha motor neurons innervate muscles collectively computes motor commands from sensory predictions, see Friston 2011: 491). Secondly, the objection assumes that structures that "merely transmit" content play no role in cognitive processing. But this is demonstrably false, for structures that do not manipulate content can play a key role in cognitive processing. For instance, one way to endow artificial neural networks with a working memory is to implement a set of self-recurring connections that do not modify the relevant activation patterns from one timestep to another (e.g. Elman

¹⁵ A critic might now point out that there are other prominent accounts of content, and that many philosophers of neuroscience endorse accounts of content based on structural similarity (e.g. Williams and Colling 2018). Surely showing that the SC satisfies teleo-informational accounts of content will not satisfy them. However, we believe that structural-similarity based accounts of content pose no challenge to our argument. This is because, as many have argued, teleo-informational accounts of content and structural similarity based ones are *not* substantially distinct; thus, everything that satisfies a teleo-informational account of content also satisfies a structural similarity based account of content (see Morgan 2014; Nirshberg and Shapiro 2020).

1991). Notice that in such a case, these connections implement a working memory precisely *because* they do not modify the relevant activation patterns across timesteps.

In conclusion, it seems that the SC easily satisfies the (parts of the) mark of the cognitive Adams and Aizawa propose. Since their proposal is the one most closely associated with a defense of contingent intracranialism, we believe that the discussion here provided is sufficient to address the contingent intracranialist worries.

4.4 - The Coupling-Constitution fallacy

Adams and Aizawa (2001; 2008; 2009; 2010; Aizawa 2010) have also repeatedly highlighted that claims of cognitive extension typically rest on what they dubbed the coupling-constitution fallacy. The fallacy consists in concluding, from the fact that cerebral processes are *causally coupled* to non-cerebral components, that cognitive processes are *partially constituted* by non-cerebral components. Just as the expansion of a bi-metallic strip of thermostat is *caused* by the environmental temperature without being constituted by it, our cognitive machinery can causally interact with the SC without being *constituted* by it.

So, have we committed the coupling-constitution fallacy? We do not think so.

Adams and Aizawa make amply clear that, in their view, putative components of cognitive systems play a constitutive (rather than merely causal) role in cognitive systems just in case they satisfy their mark of the cognitive.¹⁶ But, as we have just

¹⁶ See, for instance, (Adams and Aizawa 2010: 68): “So, if the fact that an object or process X is coupled to a cognitive agent does not entail that X is part of the cognitive agent’s cognitive apparatus, what does? The nature of X, of course. [...] One needs a theory of the ‘mark of the cognitive’”. Many thanks to an anonymous reviewer for having pointed our attention to this.

argued, the SC satisfies Adams and Aizawa's candidate mark of the cognitive. Hence, according to Adams and Aizawa, we are not committing the coupling-constitution fallacy.

But what if one does not agree with Adams and Aizawa on what tells apart causation from constitution? The literature on the extended mind is in fact rapidly converging upon a different criterion to tell apart the two: the *Mutual Manipulability* criterion (henceforth MM, see Kaplan 2012).¹⁷ We think the SC satisfies this criterion too, at least at a coarse-grained level of analysis. MM works as follows. Let ϕ be a putative component of the mechanism generating a phenomenon ψ under investigation. According to MM, ϕ is a constitutive component of the machinery producing ψ only if:

(M1) When ϕ is set to a value ϕ_1 by an (ideal) intervention, then ψ takes on a value (or some probability distribution of values) $f(\phi_1)$

(M2) When ψ is set to a value ψ_1 by an (ideal) intervention, then ϕ takes on a value (or some probability distribution of values) $f(\psi_1)$ (Craver 2007a; referenced in Kaplan 2012, p. 558).

Less formally: ϕ is a constituent of ψ only if one can experimentally manipulate ψ by manipulating ϕ (M1) and *vice versa* (M2). M1 is a *bottom-up* manipulation: the intervention focuses on a putative component, to observe whether its effect alters the phenomenon under investigation. Lesion studies are a good example. M2 speaks instead of a *top-down* manipulation: the intervention focuses on the *whole phenomenon* under investigation, to observe whether some change affects its putative components. Imaging studies are prime examples of top-down manipulations (see Craver 2007b).

¹⁷ Indeed, Gallagher (2018) reports a personal communication from Aizawa, in which the latter apparently agrees on the fact that MM is an adequate test of constitutive relevance.

Provided this minimal introduction to MM, we can now argue that the SC appears to satisfy it. Recall (Bloch *et al.* 2016) and (Vahdat *et al.* 2015). The first is a lesion study, and thus a case of bottom-up manipulation (M1). The second, instead, is a *sfMRI* experiment, and is thus a case of a top-down manipulation (M2).

Consider first the commonalities between the two experiments. In both experiments, subjects had to perform a serial reaction time task. In both cases, the subjects' task was that of pressing a sequence of buttons on a pad. And in both cases, the subjects' reaction times were measured to assess their degree of implicit learning.

Consider now the differences. In the bottom up manipulation, two groups of subjects (paraplegic subjects and healthy controls) were tested. Healthy controls exhibited the (expected) decrease of reaction times. Moreover, when a slightly different (i.e. non predictable) series of stimuli was intermixed in the trials, their reaction times increased due to an interference effect. Paraplegic subjects, in contrast, exhibited none of these two effects, suggesting that they were *not learning* the motor sequence at all. Hence, it seems that manipulating the putative component (the SC) alters the phenomenon under investigation (implicit motor learning) to a significant extent.

Consider now the *sfMRI* study of Vahdat and colleagues. Here, only healthy subjects were scanned. They were scanned in two distinct conditions. In the first, they had to press the buttons of a pad so as to act out a very simple motor sequence, which frequently recurs in normal (ecological) contexts and that they thus did not have to learn. In the second condition, they had to press the buttons acting out a fairly complex motor sequence that they most likely had to learn during the task. The

imaging data showed the presence of two spinal clusters of activity in both conditions. However, the features of these clusters differed across the two conditions. First, the amplitude of the BOLD signal was higher for the complex motor sequence. Secondly, the coordinates of the two clusters differed in the two conditions. Moreover, albeit the reaction times of subjects decreased in both conditions, only in the second condition did the rate of decrease of reaction times correlate with changes of the activity within the clusters. Lastly, only in the second condition the activity of the spinal clusters was independent from concomitant brain signals under any model of correlation investigated, reflecting an *intrinsic* contribution of spinal plasticity to the decrease of reaction times (i.e. learning). It thus seems that the manipulation of the phenomenon under investigation reverberated onto the putative component (the SC).

A similar parallelism obtains with respect to spinal involvement in affective perception. In fact, some studies presented in section 2 provide both bottom-up (lesion studies) and top-down (*sfMRI*) evidence of spinal involvement in the estimation of the affective quality of some stimuli (Pistoia *et al.* 2015; Guadagni *et al.* 2019 for bottom-up approaches; Smith and Kornelsen 2011; Kornelsen *et al.* 2015 for top-down approaches).

Is this sufficient to claim the MM criterion is satisfied? It is, but only at a *coarse grained* level of analysis. Whereas the subjects employed by Bloch and colleagues were lesioned at the level of T1-T12, Vahdat and colleagues found clusters of activity at the level of C6-C8. Nor are the tasks and stimuli of studies on emotions in subjects with SCI matched with those employed in *sfMRI* studies in healthy subjects. For this reason, we think these studies provide strong, but empirically

defeasible, evidence that the SC is a constituent of our cognitive apparatus. Yet, while not enough to *completely* defy the coupling-constitution objection, these answers seem to us sufficient to shift the burden of proof on the intracranialists' shoulders: why, in light of all this, should we consider the SC *merely* causally coupled to the cognitive system?

5 - Keeping the spinal cord in mind: why does it matter?

In this essay, we have argued that the "Cinderella view" of the SC should be abandoned, because the SC is best understood as an intrabodily cognitive extension: a piece of circuitry that, together with the brain, often and reliably conspires in producing our intelligent behaviour and generating a panoply of cognitive phenomena. Although the "Cinderella view" is the standard view of the SC in neuroscience, a triviality worry looms large over our claim.

On the one hand, Chalmers (2019) has recently argued that claims of cognitive extension which do not revolve around sensorimotor loops are just *too weak* to be interesting. On the other hand, Adams and Aizawa (2008: 17-20) suggested that considering the SC as a constituent of the cognitive apparatus would not significantly challenge their contingent intracranialism. It thus seems that both friends and foes of the extended mind will find our claim uninteresting. Yet, we believe that our claim has interesting (and unexpected) consequences, which defuse the triviality worry. We list some of them here.

To begin with, if our claim is correct, then intracranialism is false: simply put, if we are right, the cognitive system is not contained in the cranium. So we really fail to see how Adams and Aizawa can argue that considering the SC is a minor threat to

their position. More generally, if our claim is correct, then all the philosophical theories seeking to identify, reduce, or metaphysically ground by some other relation the mind in the *brain* alone are empirically inaccurate, as they cannot make justice to evidence we presented in section 2. This seems a significant change in the landscape of materialistic theories of the mind - one intracranialists should be especially wary of, given that it undercuts the claim that intracranialism is the position which best accounts for the explanatory success and the empirical practice of cognitive science (Rupert 2004; Adams and Aizawa 2008).

But perhaps intracranialists wish to defend a different claim; namely that "orthodox" cognitive science, as opposed to non-orthodox variants (e.g. Kelso 1995), supports intracranialism over any other alternative position. However, it is worth stressing that the evidence we discussed was *not* produced by researchers devoted to ecological psychology or enactive approaches to cognition. In fact, all the evidence we discussed in section 2 comes from fairly orthodox (computationalist and representationalist) approaches to cognitive neuroscience. Hence, if our claim is correct, it is simply *false* that "orthodox" approaches to cognitive science support intracranialism over alternative positions: in fact, they appear to support at least a moderate form of extracranialism.

How *moderate* is the form of extracranialism supported by "orthodox" cognitive science? For the purpose of this essay, which is squarely focused on the role of the SC in cognition, it seems to us that a satisfactory answer is this: "orthodox" cognitive science supports at least¹⁸ a strong embodied cognition thesis, according to which at

¹⁸ This "at least" qualifier is important because some authors (e.g. Rowlands 2009; Wheeler 2019) rely on experimental evidence and lines of argument largely consonant to "orthodox" cognitive science to defend extraneuralism (i.e. the claim that there are non-neural constituents of the cognitive machinery) and we do not wish to refute their claims. However, given that the focus of our essay is the role of the SC in cognition, we also do not feel obliged to provide a

least some cognitive processes are constituted by a mixture of cerebral and bodily (i.e. non-cerebral) processes (e.g. Rowlands 2010: 52-53; Newen, De Bruin and Gallagher 2018: 6). Given that the SC is not a part of the brain, spinal processes that contribute to cognition are not cerebral processes. Hence, if our claim is correct, it appears that "orthodox" cognitive science is committed to at least some strongly embodied claims. To us, this seems a surprising result, given that "orthodox" and embodied cognitive science are often pictured as mutually opposed camps (e.g. Clark 1997; Rowlands 2010).

The point just made has a further consequence regarding more moderate theoretical accounts that seek to identify a common ground between embodied and "orthodox" cognitive science by positing bodily formatted representations (e.g. Goldman 2012). According to these approaches, the body does indeed play a role in cognition, but only insofar it is represented by the brain: what counts are not real, flesh-and-blood, non-cerebral bodily parts, but just their cerebral representations (see Gallagher 2017: 28-35). If our claim is correct, however, these approaches underestimate the extent to which "orthodox" cognitive science provides an embodied view on cognition: what matters, at least in the case at hand, is not a *representation* of the SC, but the SC itself. Notice that this should not be taken as a rebuttal of bodily formatted representations: we are only claiming only that bodily formatted representations *alone* are insufficient to capture the full extent to which "orthodox" cognitive science is committed to claims about the embodiment of cognition; not that they should be abandoned.

Hence, pace Chalmers (2019) and Adams and Aizawa (2008), assigning a cognitive

defense (or an official endorsement of) extraneuralism here.

role to the SC has substantial theoretical implications, which should prompt us to critically re-evaluate what *exactly* cognitive scientists are theoretically committed to, how *exactly* classical cognitive science and 4E style cognition differ, and whether they are incommensurable or might be reconciled.

References

Adams, F. (2018). Cognition wars. *Studies in History and Philosophy of Science Part A*, 68, 20-30.

Adams F., & Aizawa K. (2001). The bounds of cognition. *Philosophical Psychology*, 14(1), 43-64.

Adams F., & Aizawa K. (2008). *The Bounds of Cognition*. Oxford: Blackwell.

Adams F., & Aizawa K. (2009). Why the mind is still in the head. In P. Robbins, & T. Aydede (Eds.), *The Cambridge Handbook of Situated Cognition* (pp. 78-95). Cambridge, MA.: Cambridge University Press.

Adams F., & Aizawa K. (2010). Defending the bounds of cognition. In R. Menary (ed.), *The Extended Mind* (pp. 67-80). Cambridge, MA.: The MIT Press.

Adams R., et al. (2013). Predictions, not commands: active inference in the motor system. *Brain Structure and Function*, 218(3), 611-643.

Ambrosini, E., et al. (2012). Which body for embodied cognition? Affordance and language within the actual and perceived reaching space. *Consciousness and Cognition*, 21(3), 1551-1557.

Aizawa K. (2010). The coupling-constitution fallacy revisited. *Cognitive System Research*, 11(4), 332-342.

Albers, A. M., et al. (2013). Shared representations for working memory and mental imagery in early visual cortex. *Current Biology*, 23(15), 1427-1431.

Alkadhi H., et al. (2005). What disconnection tells about motor imagery: evidence from parietal patients. *Cerebral Cortex*, 15(2), 131-140.

Allen, C., et al. (2009). The lower bounds of cognition: what do spinal cords reveal?. In Bickle, J. (Ed.). *The Oxford Handbook of Philosophy and Neuroscience* (pp. 129-142). New York: Oxford University Press.

Aranyosi I. (2013). *The Peripheral Mind*. New York: Oxford University Press.

Bays, P. M., & Wolpert, D. M. (2007). Computational principles of sensorimotor control that minimize uncertainty and variability. *The Journal of Physiology*, 578(2), 387-396.

Bloch A., et al. (2016). Specific deficit in implicit motor sequence learning following spinal cord injury. *PLoS ONE*, 11(6), <https://doi.org/10.1371/journal.pone.0158396>.

Boem, F., et al. (2021). Out of our skull, in our skin: the Microbiota-gut-brain axis and the extended cognition thesis. *Biology and Philosophy*, 36(2), 1-32.

Bressler, S. L., & Menon, V. (2010). Large-scale brain networks in cognition: emerging methods and principles. *Trends in Cognitive Sciences*, 14(6), 277-290.

Brozzoli, C. et al. (2011). Peripersonal space: a multisensory interface for body-object interactions. In M. M. Murray, M. T. Wallace (Eds.), *The Neural Bases of Multisensory Processes*, (pp 449-466). Boca Ranton, FL.: CRC Press.

Butt, A., & Verkhratsky, A. (2018). Neuroglia: Realising their true potential. *Brain and Neuroscience Advances*, 2. <https://doi.org/10.1177/2398212818817495>.

Chalmers D. (2008). Foreword. In A. Clark, *Supersizing the Mind* (pp. ix-xxvi). New York: Oxford University Press.

Chalmers D. (2019). Extended cognition and extended consciousness. In M. Colombo, E. Irvine, & M. Stapleton (eds.), *Andy Clark and His Critics* (pp. 9-21). New York, Oxford University Press.

Clark, A. (1997). *Being There*. Cambridge, MA.: The MIT Press.

Clark A. (2008). *Supersizing the Mind*. New York: Oxford University Press

Clark A. (2010). Memento's revenge: the extended mind, extended. In R. Menary (ed.), *The Extended Mind* (pp. 43-66). Cambridge, MA.: The MIT Press

Clark A. (2013). *Mindware: an Introduction to the Philosophy of Cognitive Science* (2nd edition). New York, Oxford University Press

Clark A. & Chalmers D. (1998). The extended mind. *Analysis*, 58(1), 7-19

Colombetti, G., & Zavala, E. (2019). Are emotional states based in the brain? A critique of affective brainocentrism from a philosophical perspective. *Biology and Philosophy*, 34(5), 1-20.

Conson M., et al. (2008). Selective motor imagery deficits in patients with locked-in syndrome. *Neuropsychologia*, 46(11), 2622-2628

Conson M., et al. (2010). Recognition and mental manipulation of body parts dissociate in locked-in syndrome. *Brain and Cognition*, 73(3), 189-193

- Costantini M., et al. (2010). Where does an object trigger an action? An investigation about affordances in space. *Experimental Brain Research*, 207(1-2), 95-103
- Costantini M., et al. (2011). When objects are close to me: affordances in the peripersonal space. *Psychonomic Bulletin & Review*, 18(2), 302-308.
- Craver C. F. (2007a). *Explaining the Brain*. New York: Oxford University Press
- Craver C. F. (2007b). Constitutive explanatory relevance. *Journal of Philosophical Research*, 32, 3-20.
- Di Pellegrino, G., & Làdavas, E. (2015). Peripersonal space in the brain. *Neuropsychologia*, 66, 126-133.
- Di Rienzo, F., et al. (2014). Impacts of neurological deficits on motor imagery: a systematic review of clinical evaluations. *Neuropsychology Review*, 24(2), 116-147.
- Dretske F. (1988). *Explaining Behavior*. Cambridge, MA.: The MIT Press.
- Elman, J. (1991). Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, 7(2-3), 195-225.
- Fiori F., et al. (2014). Motor imagery in spinal cord injury patients: moving makes the difference. *Journal of Neuropsychology*, 8(2), 199-215.
- Friston, K. (2011). What is optimal about motor control. *Neuron*, 72(3), 488-498.
- Gallagher, S. (2017). *Enactivist Interventions*. New York: Oxford University Press.
- Gallagher S. (2018). The extended mind: the state of the question. *Southern Journal of Philosophy*, 56(4), 421-447.
- Goldman A. (2012). A moderate approach to embodied cognitive science. *Review of Philosophy and Psychology*, 3(1), 71-88.
- Guadagni, V., et al (2019). Cognitive and Emotional Empathy in Individuals with Spinal Cord Injury. *Behavioural neurology*, 2019. [https://doi.org/ 10.1155/2019/1312934](https://doi.org/10.1155/2019/1312934).
- Guyton A. C., & Hall J. E. (2005). *Textbook of Medical Physiology*(11th ed.). Philadelphia: Elsevier.
- Hollerbach, J. M. (1982). Computers, brains and the control of movements. *Trends in Neurosciences*, 5, 189-192.
- Holtz A., & Levi R. (2010). *Spinal Cord Injury*. New York: Oxford University Press.
- Hubel ,D., & Wiesel ,T (1962). Receptive fields, binocular interaction, and the functional architecture of the cat's visual cortex. *The Journal of Physiology*, 160 (1),

106-154.

Hubel, D., & Wiesel, T.(1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, 19(1), 215-243.

Hurley, S. (2001). Perception and action: alternative views. *Synthese*, 129(1), 3-40.

Ionta S., et al. (2016). Spinal cord injury affects the interplay between sensorimotor and visual representations of the body. *Scientific Reports*, 6(1), 1-12.

Jeannerod M. (1994). The representing brain: neural correlates of motor intention and imagery. *Behavioral and Brain Sciences*, 17(2), 187-202.

Kandel E., et al. (2012). *Principles of Neural Science* (fifth edition). London: The McGraw-Hill Companies.

Kaplan D. M. (2012). How to demarcate the boundaries of cognition. *Biology and Philosophy*, 27(4), 545-570.

Kelso, S. (1995). *Dynamic Patterns*. Cambridge, MA.: The MIT Press.

Kirchhoff D. M. & Kiverstein J. (2019). *Predictive Processing and Extended Consciousness*. New York: Routledge.

Kornelsen, J., Smith, S. D., & McIver, T. A. (2015). A neural correlate of visceral emotional responses: evidence from fMRI of the thoracic spinal cord. *Social cognitive and affective neuroscience*, 10(4), 584-588.

Krickel B. (2019). Extended cognition, the new mechanists' mutual manipulability criterion, and the challenge of trivial extendedness. *Mind and Language*. <https://doi.org/10.1111/mila.12262>

Logothetis, N. K. (2002). The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1424), 1003-1037.

Ludwig, D. (2015). Extended cognition and the explosion of knowledge. *Philosophical Psychology*, 28(3), 355-368.

Lungu O., et al. (2010). Changes in spinal reflex excitability associated with motor sequence learning. *Journal of Neuropsychology*, 103(5), 2675-2683.

Menary R. (ed.). (2010). *The Extended Mind*. Cambridge, MA.: The MIT Press.

McCauley, R. N., & Bechtel, W. (2001). Explanatory pluralism and heuristic identity theory. *Theory and Psychology*, 11(6), 736-760.

Miller, K. J., et al. (2010). Cortical activity during motor execution, motor imagery,

and imagery-based offline feedback. *Proceedings of the National Academy of Sciences*, 107(9), 4430-4435.

Montoya, P., & Schandry, R. (1994). Emotional experience and heartbeat perception in patients with spinal cord injury and control subjects. *Journal of Psychophysiology*, 8, 289-289.

Morgan, A. (2014). Representations gone mental. *Synthese*, 191(2), 213-244.

Neander, K. (2017). *A Mark of the Mental*. Cambridge, MA.: The MIT Press.

Newen A., De Bruin L., Gallagher S. (2018). 4E cognition: historical roots, key concepts and central issues, in A. Newen, L. De Bruin, S. Gallagher (eds.), *The Oxford Handbook of 4E Cognition* (pp. 3-19). New York: Oxford University Press.

Nirshberg, G., & Shapiro, L. (2020). Structural and Indicator representations: a difference in degree, not in kind. *Synthese*, <https://doi.org/10.1007/s11229-020-02537-y>.

Olsson C. J. (2012). Complex motor representations might not be preserved after spinal cord injury. *Experimental Neurology*, 236(1), 46-49.

Parsons L. M. (1987). Imagined spatial transformations of one's hand and feet. *Cognitive Psychology*, 19(2), 178-241.

Parsons L. M. (1994). Temporal and kinematic properties of motor behavior reflected in mentally simulated action. *Journal of Experimental Psychology: Human Perception and Performance*, 20(4), 709-730.

Pistoia, F., et al. (2015). Contribution of interoceptive information to emotional processing: evidence from individuals with spinal cord injury. *Journal of neurotrauma*, 32(24), 1981-1986.

Ramsey, W. (2017). Must cognition be representational? *Synthese*, 194(11), 4197-4214.

Rizzolatti G., et al. (1997). The space around us. *Science*, 277(5323), 190-191.

Robertson E. M. (2007). The serial reaction time task: implicit motor skill learning?. *Journal of Neuroscience*, 27(38), 10073-10075.

Rowlands, M. (2009). Extended cognition and the mark of the cognitive. *Philosophical Psychology*, 22(1), 1-19.

Rowlands M. (2010). *The New Science of the Mind*. Cambridge, MA.: The MIT Press.

Rupert R. (2004). Challenges to the hypothesis of extended cognition. *The Journal of Philosophy*, 101(8), 389-428.

Rupert R. (2009). *Cognitive Systems and the Extended Mind*. New York: Oxford University Press.

Rupert R. (2010). Extended cognition and the priority of cognitive systems. *Cognitive Systems Research*, 11(4), 343-356.

Scandola M., et al. (2016). Spinal cord lesions shrink peripersonal space around the feet, passive mobilization of the paraplegic limb restores it. *Scientific Reports*, 6, 24126.

Scandola, M., et al. (2020). Visuo-motor and interoceptive influences on peripersonal space representation following spinal cord injury. *Scientific Reports*, 10(1), 1-16.

Sedda A., et al. (2019). Affordances after spinal cord injury. *Journal of Neuropsychology*, 13(2), 354-369.

Shea, N. (2018). *Representation in Cognitive Science*. New York: Oxford University Press.

Smith, S. D., & Kornelsen, J. (2011). Emotion-dependent responses in spinal cord neurons: a spinal fMRI study. *NeuroImage*, 58(1), 269-274.

Smith, S. D., Kolesar, T. A., & Kornelsen, J. (2018a). Lateralized brainstem and cervical spinal cord responses to aversive sounds: A spinal fMRI study. *Brain sciences*, 8(9), 165.

Smith, S. D., Kornelsen, J., & McIver, T. A. (2018b). Generating facial expressions of disgust activates neurons in the thoracic spinal cord: a spinal fMRI study. *Social neuroscience*, 13(3), 328-332.

Sprevak M. (2009). Extended cognition and functionalism. *The Journal of Philosophy*, 106(9), 503-527.

Stroman, P.W., et al. (2014). The current state-of-the-art of spinal cord imaging: methods. *Neuroimage* 84, 1070-1081.

Tao, D, et al. (2021). Where does fear originate in the brain? A coordinate-based meta-analysis of explicit and implicit fear processing. *NeuroImage*, 227, 117686.

Todorov E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9), 907-915.

Vahdat S., et al. (2015). Simultaneous brain-cervical cord fMRI reveals intrinsic spinal cord plasticity during motor sequence learning. *PLoS Biology*, 13(6). <https://doi.org/10.1371/journal.pbio.1002186>.

Wheeler, M. (2011). In search of clarity about parity. *Philosophical Studies*, 152(3),

417-425.

Wheeler M. (2019). Breaking the waves. In M. Colombo, E. Irvine, & M. Stapleton (eds.), *Andy Clark and His Critics* (pp. 81-95). New York, Oxford University Press.

Wheeler-Kingshott, C.A., et al. (2013). The current state-of-the-art of spinal cord imaging: applications. *Neuroimage* 84, 1082-1093.

Williams, D., & Colling, L. (2018). From symbols to icons: the return of resemblance in the cognitive neuroscience revolution. *Synthese*, 194(5), 1941-1967.

Wolpaw J. R. (2007). Spinal cord plasticity in acquisition and maintenance of motor skills. *Acta Psychologica*, 189(2), 155-169.

Wolpaw J. R., & Tennissen A. M. (2001). Activity-dependent spinal cord plasticity in health and disease. *Annual Review of Neuroscience*, 24(1), 807-843.