

Modelling aspects of consciousness: a topological perspective

Mike Steel

Biomathematics Research Centre, University of Canterbury, Christchurch, New Zealand

Abstract

Attention Schema Theory (AST) is a recent proposal to provide a scientific explanation for the basis of subjective awareness. In AST, the brain constructs a representation of attention taking place in its own (and others') mind ('the attention schema'). Moreover, this representation is incomplete for efficiency reasons. This inherent incompleteness of the attention schema results in the inability of humans to understand how their own subjective awareness arises (related to the so-called 'hard problem' of consciousness). Given this theory, the present paper asks whether a mind (either human or machine-based) that incorporates attention, and that contains a representation of its own attention, can ever have a complete representation. Using a simple yet general model and a mathematical argument based on classical topology, we show that a complete representation of attention is not possible, since it cannot faithfully represent streams of attention. In this way, the study supports one of the core aspects of AST, that the brain's representation of its own attention is necessarily incomplete.

Email: mike.steel@canterbury.ac.nz

Keywords: consciousness, neuroscience, attention schema theory, topology, dimensionality

1. Introduction

The question of explaining the basis and nature of human consciousness remains an outstanding scientific question that straddles several disciplines, from neuroscience, cognitive psychology and biophysics (Baars, 1988; Dehaene et al., 2017; Lau and Rosenthal, 2011; Llinás et al., 1998; Rolls, 2007) to philosophy (Chalmers, 1995; Dennett, 1992). The question is often divided into two parts: the so-called ‘easy’ and ‘hard’ problems (Chalmers, 1995). Roughly speaking, the ‘easy’ problem seeks to identify the neural processes that give rise to the experiences one typically associates with consciousness (perception, subjective awareness, introspection, thought and emotion). The second (‘hard’) problem is more fundamental: it asks how is it that humans experience subjective awareness at all (i.e. the experience of ‘qualia’), rather than simply processing information in an unconscious but purely computer-like fashion.

A great deal has been learned in recent decades concerning the first problem, and theories have been proposed to describe how conscious experience might arise within a brain or within a sufficiently complex neural network. Two leading approaches are (neuronal) Global Workspace Theory (GWT) (Baars, 1988; Wallace, 2005) and Integrated Information Theory (Esteban et al., 2018; Tononi, 2008), though many other related ideas have been put forward, including, for example, thalamocortical resonance (Llinás et al., 1998) or the emergence of consciousness as a phase transition in brain dynamics (Werner, 2013).

Here, we consider an approach to consciousness that deals with both problems, namely Attention Schema Theory (AST), pioneered by Michael Graziano and colleagues (Graziano and Webb, 2014; Graziano, 2014, 2019; Webb and Graziano, 2015). AST builds on approaches such as GWT to describe how attention arises within a brain by well-studied processes in which the focus of attention (which can be on external stimuli or internally on memories or thoughts) is governed by a competitive process of excitation, reinforcement and inhibition. AST further suggests that conscious awareness requires an additional aspect – an internal concept of the self, and an internal cognitive model of how attention works (the ‘Attention Schema’) – the latter having evolved by natural selection in higher animals, because of the selective advantage of the brain controlling its own attention processes, as well as providing a way for an individual to model the minds of other individuals in social settings (Graziano and Webb, 2014).

According to AST, a central reason why the ‘hard’ problem of consciousness appears hard is simply because the attention schema is highly incomplete, because it is a grossly oversimplified internal model that does not include all the neural details of how attention actually arises (greater completeness and additional neural detail in an attention schema would reduce computational efficiency, and so be unlikely to evolve). This incompleteness leads individuals to conclude they have a mysterious essence inside of them, consciousness, instead of the physically real process of attention that they actually have.

Thus, AST provides a specific structural process for consciousness to arise and for a mind to know that it is conscious, yet without having the ability to completely understand why. AST also suggests that current machines are unlikely to be conscious but could be designed to be so (though it is not enough to simply increase computational power). For further details, the reader is referred to Graziano and Webb (2014); Graziano (2014, 2019); Webb and Graziano (2015).

In this paper, we provide a simple mathematical analysis of a question that arises from AST. We ask whether a mind could, in principle, have an attention schema that is complete, in the sense that any stream of attention – at least over a short time-frame – can be faithfully represented within the attention schema. Using a simple argument based on standard concepts from topology, we show that a complete attention schema is inherently impossible, providing a further and more fundamental reason for the incompleteness assumption in AST.

2. Mathematical model

The model we describe is particularly simplified; however, by making minimal assumptions and working in a high-level setting, the results we describe — while less detailed — are more generic than those that might be derived from a detailed and finely-tuned model. Our approach is based on concepts and results from topology, which is the mathematical study of ‘space’ (and maps between spaces) in a very general sense (topological spaces include many familiar examples that form the basis of physics, biology and other fields in science) (Ghrist, 2014; Hatcher, 2002).

We view attention within an individual mind as moving across some open, bounded and connected subspace \mathcal{A} of a high-dimensional Euclidean space (thus \mathcal{A} comprises a n -dimensional manifold, with each point in \mathcal{A} having an

open neighbourhood that is (homeomorphic to) an open n -cell). The space \mathcal{A} could, for example, describe neuronal firing patterns, fluctuating neural states and changing parameters within the brain; alternatively, \mathcal{A} might be taken to be a purely phenomenological space, or perhaps some space that involves aspects of both types. In this paper, we impose no restriction on n ; for example, n could be in the order of the number of synaptic junctions in a brain or greater.

Throughout we will adopt a convention in topology of using the word *map* to refer to any continuous function; also, a map is said to be *injective* if it is one-to-one.

We will let o_t denote an open set in \mathcal{A} that is the focus of attention at time t (thus, o_t is also a manifold of dimension n). Consistent with approaches such as GWT, we will describe o_t as the region of \mathcal{A} that has reached the threshold required for attention at the given instant t . More formally, consider a function F that assigns a non-negative real value to each pair (t, x) (t =time, x is an element of \mathcal{A}); we will assume throughout that F is continuous in both variables and refer to it as an *attentional activity function*.

Given a threshold $\delta > 0$, we will formally define o_t to be the portion of \mathcal{A} that lies above this threshold at time t and thereby becomes the focus of attention at time t .¹ More precisely:

$$o_t = \{x \in \mathcal{A} : F(t, x) > \delta\}. \quad (1)$$

Notice that F may vary with t while o_t remains constant.²

We now introduce two further (related) definitions. First, given an interval I of strictly positive length, we say that o_t is a *frozen stream of attention* if, for all $t, t' \in I$ with $t < t'$, we have $o_t = o_{t'}$. Thus a frozen stream of attention is one where the focus of attention does not shift (as in Fig. 1 between t_0 and t_1). Next, given an (open or closed) interval I , let O_I denote the subspace of $I \times \mathcal{A}$ defined by:

$$O_I = \{(t, x) : t \in I, x \in o_t\},$$

¹One could further allow δ to depend continuously on t , but this gives no further generality since δ can be absorbed into F ; indeed, by rescaling, one could assume that $\delta = 1$.

²In topological terms, o_t need not be connected (we give an example shortly), or even contractible, and o_t could also contain ‘holes’ of varying dimensions (i.e. non-trivial homology).

and consider the quotient space:

$$o_I = O_I / \sim,$$

where \sim is the equivalence relation on O_I defined by $(t, x) \sim (t', x')$ if $x = x'$ and o_t is a frozen stream of attention on the closed interval with endpoints t and t' . Thus, O_I describes the stream of attention over the time period I , whereas the quotient space o_I collapses any portions of this space that consist of frozen streams of attention (thus, o_I describes just the changes in attention over the time interval I). The motivation for introducing this quotient step is to allow the possibility for a mind to hold a representation of attention over time if the attention is not shifting with time, and thereby strengthening the statement of our main result.

2.1. Attention schema

A *formal attention schema*³ for (\mathcal{A}, F) is a connected open subspace \mathcal{A}_S of \mathcal{A} together with a map φ_I from any quotient (fixed or moving) stream of attention o_I for \mathcal{A} into \mathcal{A}_S . We denote the collection of the maps φ_I (over all intervals I and for a fixed \mathcal{A}) by φ .

A simple but trivial attention schema is provided by selecting some fixed point $a \in \mathcal{A}_S$, and letting φ map each point in o_I (for all I) to a in \mathcal{A}_S . A slightly less trivial example, would be to map each o_I to its own particular (distinct) point in \mathcal{A}_S ; however, this still misses much of the detail of o_I , even when o_t is frozen in the interval I . To provide a formal attention schema for (\mathcal{A}, F) that has greater representational detail we will say that (\mathcal{A}_S, φ) is *complete* if φ_I is injective for all time intervals I of length up to some given (positive) duration T , otherwise, we say that (\mathcal{A}_S, φ) is *incomplete* (the latter means that different streams of attention (up to duration T) cannot be distinguished, as they map to the same state in the attention schema).

Note that we are not necessarily assuming (in a given mind) that each stream of attention o_I is continually being encoded in \mathcal{A}_S . Rather, completeness simply means that the attention schema is sufficiently detailed that it has the capability to do so over any (short but positive-length) interval, and in such a way that the encoding is faithful (i.e. injective).

³The word ‘formal’ is to emphasise a mathematical abstraction of the notion from AST.




Figure 1: A schematic illustration of the model in a highly simplified setting where o_t is two-dimensional. *Left*: A stream of attention o_t of an individual mind described by the portion of \mathcal{A} that has attentional activity (indicated by F) larger than a threshold (δ). Initially, attention is focused on the apple but then shifts when a mouse appears, and eventually eats the apple. The open sets o_t are two-dimensional (in the horizontal plane) and involve a frozen stream of attention on the interval $[t_0, t_1]$, after which a jump occurs (when the mouse is noticed), resulting in o_t briefly having two components. Subsequently, o_t contains two (or more) moving streams of attention as the mouse moves towards the apple and then begins to eat it. *Middle*: The streams of attention in O_I map to the quotient space o_I (i.e. the initial frozen stream of attention has been collapsed). *Right*: A mapping of o_I into a formal attention schema \mathcal{A}_S . Since \mathcal{A}_S has the same dimension as \mathcal{A} in this example and o_t contains a moving stream of attention, the map cannot be injective (by Proposition 1(i)).

2.2. Moving streams of attention

In contrast to a frozen stream of attention, we say that o_t as defined by (1) is a *moving stream of attention*, as t varies over an interval I , provided that the following two conditions both hold for each $t \in I$:

- (m-i) If $o_t = o_{t'}$ for $t' \in I$, then $t = t'$.
- (m-ii) o_t can be written as $\{g(t, x) : x \in o\}$, where o is an open set in \mathcal{A} and $g : I \times o \rightarrow \mathcal{A}$ is a map for which $x \mapsto g(t, x)$ is injective.

In words, a moving stream of attention is one that moves continuously over a period of time (e.g. across \mathcal{A} , or by becoming larger or smaller) without either stalling or returning to a previous state (i.e. violating (m-i)) or suddenly

jumping to a different region of \mathcal{A} (i.e. violating (m-ii)). Fig. 1 illustrates these notions in a simple and low-dimensional setting.

An example of a moving stream of attention would be the continuous changes in attention from observing the mouse move towards the apple in Fig. 1. Notice that although o_t in Fig. 1 contains moving streams of attention, o_t is not a moving stream of attention over the entire interval from t_2 to t_3 . To see why, observe that Condition (m-ii) implies that o_t and $o_{t'}$ are homeomorphic (i.e. topologically equivalent) for any pair $t, t' \in I$, since they are both isotopic embeddings of o into \mathcal{A} ; in particular, o_t is connected if and only if $o_{t'}$ is also connected (however, o_t is disconnected at t_2 and connected at t_3).

It may be tempting to suppose that if O_I is not a frozen stream of attention, then it must contain a moving stream of attention on some sub-interval I' of strictly positive length. However, it is easy to construct counterexamples, even when o_t is described by Eqn. (1), as we will formally demonstrate in the next section.

3. Formal result

In the following result, recall that an injective map refers to any continuous function that is one-to-one (i.e. no two elements are mapped to the same point).

Proposition 1. *Suppose that (\mathcal{A}_S, φ) is a formal attention schema for (\mathcal{A}, F) and I is an interval of time that has positive length. The following then hold.*

- (i) *If O_I contains a moving stream of attention (on a subinterval of I that has strictly positive length), then o_I cannot be mapped injectively into \mathcal{A}_S . In particular, a complete attention schema is not possible in the presence of moving streams of attention.*
- (ii) *For any sequence o_1, o_2, \dots, o_k of open sets in \mathcal{A} (with no bound on k), there is a (continuous) attentional activity function $F = F(t, x)$ for which O_I contains consecutive frozen streams of attention consisting of o_1, o_2, \dots , and for which o_I can be mapped injectively into \mathcal{A}_S .*

The proof of Proposition 1 is given in the Appendix. Here, we make two remarks. Firstly, although Part (ii) suggests a way in which a stream of attention is moving over the interval I , the changes are not continuous; instead, they proceed via a sequence of ‘jumps’ (each of which may be small, though some (or all) could also be large). Second, if we were to allow part or all of the formal attention schema \mathcal{A}_S to inhabit a space of higher dimension than the rest of \mathcal{A} (i.e. to be an open bounded subspace of \mathbb{R}^N for $N > n$), then moving streams of attention (not involving \mathcal{A}_S) could now be mapped injectively into \mathcal{A}_S . However, this does still not allow for a complete formal attention schema, since attention could now move across this higher-dimensional part of \mathcal{A}_S (i.e. a moving stream of attention across the attention schema itself).

4. Concluding comments

Proposition 1 required a number of assumptions, and we comment briefly on the impact of relaxing or removing them. Firstly, if \mathcal{A} is allowed to have cells of infinite or arbitrarily high dimension (e.g. to be a subspace of the countably infinite union of the finite-dimensional spaces \mathbb{R}^i for $i = 1, 2, 3 \dots$), then Proposition 1(i) need not hold. Similarly, if \mathcal{A} is allowed to vary with t , so that cells of increasing dimension can continue to be generated as time advances, then Proposition 1(i) may again not hold (at any given time). However, if one accepts that a mind is entirely dependent on a nervous system (and external stimuli), which, though extremely complex, are nevertheless of bounded dimensions, the assumption of a finite (albeit large) dimensional topological space seems reasonable.

Our theoretical result offers some insight into the generic properties of such models such as AST. Other mathematical approaches to consciousness have been explored in a number of recent papers (see e.g. Aleksander (2020); Maguire et al. (2016); Moyal et al. (2020); Rudrauf et al. (2017); Signorelli and Meling (2021); Velazquez (2020)); however, the approach here is quite different. The use of topology has proven useful in identifying generic properties for a range of other processes and systems in applied science (see e.g. Ghrist (2014)). Indeed, an applied branch of algebraic topology — persistent homology — has led to some important data-driven insights into the structural topology of brain processes, including the neural correlates of consciousness (Expert et al., 2019; Petri et al., 2014); again, the approach here (which focuses on AST) has no direct relationship to this earlier work.

5. Declaration of Competing Interest

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

6. Acknowledgements

I wish to thank the reviewers for several helpful comments on an earlier version of this manuscript, and Jasmine Peate-Garratt for sketching some images used in Fig. 1.

References

- Aleksander, I., 2020. The category of machines that become conscious. *Journal of Artificial Intelligence and Consciousness* 7, 3–13.
- Baars, B.J., 1988. *A Cognitive Theory of Consciousness*. Cambridge Univ. Press.
- Chalmers, D.J., 1995. Facing up to the problem of consciousness. *Journal of Consciousness Studies* 2, 200–219.
- Dehaene, S., Lau, H., Kouider, S., 2017. What is consciousness, and could machines have it? *Science* 358, 486–492.
- Dennett, D.C., 1992. *Consciousness Explained*. Back Bay Books, New York.
- Esteban, F.J., Galadí, J.A., Langa, J.A., Portillo, J.R., Soler-Toscano, F., 2018. Informational structures: A dynamical system approach for integrated information. *PLoS Computational Biology* 14, e1006154. DOI: 10.1371/journal.pcbi.1006154.
- Expert, P., Lord, L.D., Kringelbach, M.L., Petri, G., 2019. Editorial: Topological neuroscience. *Networks Neuroscience* 3, 653–655.
- Ghrist, R., 2014. *Elementary Applied Topology*. CreateSpace Independent Publishing Platform, USA.
- Graziano, M.S.A., 2014. Speculations on the evolution of awareness. *Journal of Cognitive Neuroscience* 26, 1300–1304. DOI: 10.1162/jocn_a_00623.

- Graziano, M.S.A., 2019. *Rethinking Consciousness: A Scientific Theory of Subjective Experience*. W. W. Norton and Co., New York, NY.
- Graziano, M.S.A., Webb, T.W., 2014. A mechanistic theory of consciousness. *International Journal of Machine Consciousness* 6, 163–176.
- Hatcher, A., 2002. *Algebraic Topology*. Cambridge University Press, NY, USA. ISBN: 9780521795401.
- Lau, H., Rosenthal, D., 2011. Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Science* 15, 365–373. DOI: 10.1016/j.tics.2011.05.009.
- Llinás, R., Ribary, U., Contreras, D., Pedroarena, C., 1998. The neuronal basis for consciousness. *Philosophical Transactions of the Royal Society of London B (Biological Sciences)* 353, 1841–1849. DOI: 10.1098/rstb.1998.0336.
- Maguire, P., Moser, P., Maguire, R., 2016. Understanding consciousness as data compression. *Journal of Cognitive Science* 17, 63–94.
- Moyal, R., Fekete, T., Edelman, S., 2020. Dynamical emergence theory (DET): A computational account of phenomenal consciousness. *Minds and Machines* 30, 1–21. DOI: 10.1007/s11023-020-09516-9.
- Petri, G., Expert, P., Turkheimer, F., et al., 2014. Homological scaffolds of brain functional networks. *Journal of the Royal Society Interface* 11, 20140873. DOI: 10.1098/rsif.2014.0873.
- Rolls, E.T., 2007. A computational neuroscience approach to consciousness. *Neural Networks* 20, 962–982.
- Rudrauf, D., Bennequin, D., Granic, I., Landini, G., Friston, K., Williford, K., 2017. A mathematical model of embodied consciousness. *Journal of Theoretical Biology* 428, 106–131.
- Signorelli, C.M., Meling, D., 2021. Towards new concepts for a biological neuroscience of consciousness. *Cognitive Neurodynamics* (in press).
- Tononi, G., 2008. Consciousness as integrated information: a provisional manifesto. *Biological Bulletin* 215, 216–242. DOI: 10.2307/25470707.

- Velazquez, J.L.P., 2020. On the emergence of cognition: from catalytic closure to neuroglial closure. *Journal of Consciousness Studies* 46, 95–119.
- Wallace, R., 2005. *Consciousness: A Mathematical Treatment of the Global Neuronal Workspace Model*. Springer, USA.
- Webb, T.W., Graziano, M.S.A., 2015. The attention schema theory: A mechanistic account of subjective awareness. *Frontiers in Psychology* 6, 500. DOI: 10.3389/fpsyg.2015.00500.
- Werner, G., 2013. Consciousness viewed in the framework of brain phase space dynamics, criticality, and the renormalization group. *Chaos, Solitons and Fractals* 55, 3–12.

Appendix: Proof of Proposition 1

Part (i): Suppose that $O_{I'}$ is a moving stream of attention (for some subinterval I' of I of strictly positive length). Recall that $O_{I'} = \{(t, x) : t \in I', x \in o_t\} \subseteq I' \times \mathcal{A}$, and $o_{I'}$ is the quotient space $O_{I'}/\sim$. By Condition (m-i), no two points of $O_{I'}$ are equivalent under the relation \sim , so $o_{I'}$ is homeomorphic to $O_{I'}$.

Now consider the function g that is the subject of Condition (m-ii) (using I' in place of I), let

$$Y = \{(t, x) : t \in I', x \in o\},$$

and let

$$H : Y \rightarrow O_{I'}$$

be defined by $H(t, x) = (t, g(t, x))$.

We claim that H is an injective map. To see this, observe that the continuity of H follows from the continuity of g , whereas for injectivity, the equality $H(t, x) = H(t', x')$ implies that $(t, g(t, x)) = (t', g(t', x'))$ and so $t = t'$, and hence $g(t, x) = g(t, x')$, thus $x = x'$ by the injectivity of the function $x \mapsto g(t, x)$ for each fixed t , from Condition (m-ii).

The set Y contains an open cell e_{n+1} of dimension $n + 1$ (since I' contains an open 1-cell, and o contains an open n -cell, and the resulting product space of these two cells is an open $(n + 1)$ -cell).

Let us now suppose that φ_I is injective (we will show that this is impossible by deriving a contradiction). We then have the following composition of injective maps:

$$e_{n+1} \xrightarrow{i} Y \xrightarrow{H} O_{I'} \xrightarrow{\epsilon} o_I \xrightarrow{\varphi_I} \mathcal{A}_S \xrightarrow{i'} \mathbb{R}^n, \quad (2)$$

where i and i' are subspace inclusion maps, and $\epsilon : O_{I'} \rightarrow o_I$ maps x to $[x]$ where $[*]$ refers to the equivalence class under \sim (ϵ is injective by Condition (m-i)).

We now apply a classical result from algebraic topology: the Invariance of Domain theorem (e.g. Hatcher (2002)). One form of this theorem states that if $f : o \rightarrow o'$ is an injective map from an open m -cell to \mathbb{R}^n , then $m \leq n$. Applying this with $m = n + 1$, and noting that the composition of injective maps in (2) gives an injective map $f = i' \circ \varphi_I \circ \epsilon \circ H \circ i$ from e_{n+1} into \mathbb{R}^n which results in the contradiction $n + 1 \leq n$. Thus φ_I cannot be injective, as claimed.

Part (ii): We first show that if o is an open set in \mathcal{A} and $\delta > 0$ is a threshold, then there is a continuous function $\psi_o : \mathcal{A} \rightarrow \mathbb{R}^{\geq 0}$ for which $o = \{x \in \mathcal{A} : \psi_o(x) > 0\}$. To see why this holds, let $\partial(o)$ denote the boundary (topological frontier) of o in \mathcal{A} (since o is open in \mathcal{A} , this is the set of points in \mathcal{A} that are in the closure of o but not in o), and let $D(x, \partial(o)) = \inf\{d(x, y) : y \in \partial(o)\}$, where d is any standard metric on \mathbb{R}^n (e.g. the Euclidean metric). Since $\partial(o)$ is closed and bounded, and therefore compact (by the Heine-Borel theorem), it follows that $D(x, \partial(o))$ is realised by at least one point $y \in \partial(o)$. Let:

$$\psi_o(x) = \begin{cases} D(x, \partial(o)), & \text{if } x \in o; \\ 0, & \text{if } x \in \mathcal{A} - o. \end{cases}$$

The function ψ_o is well-defined, continuous, and $\psi_o(x) > 0$ if $x \in o$ and $\psi_o(x) = 0$ if $x \in \mathcal{A} - o$. In particular, $o = \{x \in \mathcal{A} : \psi_o(x) > 0\}$, as claimed.

Without loss of generality, we may assume that I is the open interval $(0, 1)$. Write I as a sequence of open intervals of I of equal length interspersed with singleton elements of I as follows:

$$(0, \frac{1}{k}), [\frac{1}{k}], (\frac{1}{k}, \frac{2}{k}), [\frac{2}{k}], (\frac{2}{k}, \frac{3}{k}) \dots [\frac{k-1}{k}], (\frac{k-1}{k}, 1).$$

For $j = 1, \dots, k$, let I_j denote the open interval $(\frac{j-1}{k}, \frac{j}{k})$, and for $t \in I_j$ let

$$\gamma_j(t) = \begin{cases} \min \{t - \frac{j-1}{k}, \frac{j}{k} - t\}, & \text{if } t \in I_j; \\ 0, & \text{otherwise.} \end{cases}$$

Thus $\gamma_j(t)$ is continuous on I , and is non-zero and strictly positive precisely on I_j . We now define an attentional activity function $F : (0, 1) \times \mathcal{A} \rightarrow \mathbb{R}^{\geq 0}$ by setting:

$$F(t, x) = \delta + \sum_{j=1}^k \gamma_j(t) \cdot \psi_{o_j}(x).$$

The function F is then well-defined, continuous in both variables, and F has the property that o_t defined by Eqn. (1) coincides with o_j throughout I_j for $j = 1, \dots, k$ (at the singleton element $t_j = \frac{j}{k}$ that appears between I_j and I_{j+1} for $j = 1, \dots, k-1$ we have: $o_{t_j} = \emptyset$). Thus O_I contains frozen versions of o_1, o_2, \dots, o_k (in this order) and o_I consists of k connected components, namely disjoint copies of o_1, o_2, \dots, o_k . Since each set o_j is an open set in \mathcal{A} (and therefore an open set of \mathbb{R}^n), there is an injective map of each set o_j into any given open n -cell e_n contained in \mathcal{A}_S ; therefore, there is an injective map of o_I into e_n as well (by selecting k disjoint open n -cells within e_n to map each o_j set into).

Note that our construction provides just one way of guaranteeing the existence of a (continuous) attentional activity function F satisfying the property required; many other possible choices for F are possible. \square