



Mapping the philosophy and neuroscience nexus through citation analysis

Eugenio Petrovich¹ · Marco Viola²

Received: 11 October 2023 / Accepted: 5 November 2024
© The Author(s) 2024

Abstract

We provide a quantitative analysis of the philosophy-neuroscience nexus using citation analysis. Combining bibliometric indicators of cross-field visibility with journal citation mapping techniques, we investigate four dimensions of the nexus: how the visibility of neuroscience in philosophy and of philosophy in neuroscience has changed over time, which areas of philosophy are more interested in neuroscience, which areas of neuroscience are more interested in philosophy, and how the trading zone between the two fields is configured. We also discuss two hypotheses: the supposed occurrence of a neuro-revolution in philosophy and the role of psychology as the disciplinary link between neuroscience and philosophy. Both the visibility of neuroscience in philosophy and the visibility of philosophy in neuroscience have increased significantly from 1980 to 2020, albeit the latter remains an order of magnitude lower than the former. Neuroscience is particularly visible in philosophy of mind, applied ethics, philosophy of science, but not in ‘core’ areas of analytic philosophy. Philosophy is particularly visible in cognitive and systems neuroscience and neuropsychiatry, but not in biomedical neuroscience. As for the trading zone between philosophy and neuroscience, our data show that it works differently in philosophy and in neuroscience. While some philosophy journals are active loci of bidirectional communication, neuroscience journals are divided between journals ‘importing’ philosophy in neuroscience and journals ‘exporting’ neuroscience to philosophy. Lastly, data do not support the hypothesis that a widespread neuro-revolution has transformed philosophy radically, but support the hypothesis that psychology functions as a mediating disciplinary link between philosophy and neuroscience.

Keywords Philosophy and neuroscience · Citation analysis · Bibliometrics · Science mapping

✉ Eugenio Petrovich
eugenio.petrovich@unito.it

¹ University of Turin, Department of Philosophy and Education Sciences, Torino, Italy

² University of Roma Tre, Department of Philosophy, Communication, and Performing Arts, Rome, Italy

1 Introduction

In recent decades, there has been a growing connection between philosophy and neuroscience. A glance at the contemporary *milieu* of philosophy reveals a sheer number of conferences, publications, and institutionalized activities focused on neuroscience. Subfields such as neuroethics have settled their own scientific societies and journals, several forums have been established for promoting the dialogue between philosophers and neuroscientists, and new journals have been founded for interdisciplinary research.

Still, while some quantitative studies have been recently performed on some specific points of contact between philosophy and neuroscience (Kostić & Halffman, 2023; Leefmann et al., 2016; Yan & Liao, 2023), no *global* quantitative analysis of the interaction between philosophy and neuroscience exists. With the current study, we aim to address this gap using citation analysis, which has been recently successfully applied to the investigation of other science-philosophy connections (Bonino et al., 2022; Khelifaoui et al., 2021; McLevey et al., 2018; Pradeu et al., 2021).

The paper is structured as follows. In the next section, we provide a qualitative reconstruction of the history of interactions between philosophy and neuroscience over the last forty years, based on the perspectives of key figures in the field. In the subsequent section, we formulate four research questions and two hypotheses concerning the characteristics of the philosophy-neuroscience nexus. The Methodology section justifies the use of citation analysis to address these questions and explains in detail the bibliometric operationalizations and indicators we developed. Our main results are presented in the Results section and interpreted in light of the four questions and two hypotheses in the Discussion section. Lastly, the Conclusion section reviews some limitations of the present study and proposes further lines of inquiry for future research.

While a linear reading is advised in order to fully understand the nature and rationale of our study, an impatient reader should be able to skip the Methodology and the Result sections altogether, only glancing at the figures therein, and still understand our overarching narrative.

2 Philosophers' contact with neuroscience

Philosophy and neuroscience are traditionally enlisted among the six disciplines leading the cognitive revolution in the Sixties, together with linguistics, artificial intelligence, anthropology, and psychology (Gardner, 1987). Nonetheless, by the end of the Seventies, the well-known *Report to the Sloan Foundation on the status of cognitive science* claimed that, unlike what happened in the other pairs of disciplines, the topics at the intersection between philosophy and neuroscience “have not yet become the focus of formally recognized scholarly effort” (State of the Art Committee, 1978).

However, a turnabout was right behind the corner. Or at least, that is what Patricia Churchland's influential call to arms was hoped to achieve. Starting with her manifesto book *Neurophilosophy* (1986), she invoked a "paradigm shift" (Churchland, 1987, p. 545) in the relationship between philosophy and the neuroscience. Loyal to the Quinean idea that scientific and philosophical problems lie on a continuum, Churchland argued that philosophical debates should pay more attention to the latest discoveries from the neuroscience and less to old methods such as linguistic analysis.

The Nineties afforded a fertile ground for cultivating philosophers' interest in neuroscience, not only because of the general neuro-hype that fueled big science initiatives like "the decade of the brain" (1990–1999), which was accompanied by an increase of neuroscientific papers (see Fig. 2C below), but also because the development of neuroimaging techniques brought neuroscience—a field unfamiliar to most philosophers at the time—closer to psychology, a discipline with which philosophers were more familiar (Cooper & Shallice, 2010). As a result, by the end of the century, philosophy and neuroscience got increasingly closer. In 2000, John Bickle – one of the most enthusiastic adherents of Churchland's neurophilosophy – founded the journal *Brain and Mind*, to establish "a forum for neuroscientists and neurophilosophers to discuss larger interdisciplinary, scientific, methodological, philosophical, and social consequences of contemporary neuroscience" (Bickle et al., 2000, p. 2).

Despite the aspirations of the most fervent "neurophilosophers", however, the journal ceased publication after only four years due to a dearth of submissions and was subsequently incorporated into *Synthese* (Bickle, 2003a). Apparently, what Bickle and many others took to be a Kuhnian paradigm-shift, warranting dedicated publication outlets, turned into a reform and found its niche in the existing editorial landscape (Bickle, 2019).

In spite of the retreat from the revolutionary spirit, the new millennium witnessed the increasing institutionalization of philosophers' interest in neuroscience. The entry "The Philosophy of Neuroscience" was published on the Stanford Encyclopedia of Philosophy in 1999, and regularly updated (Bickle et al., 1999/2019); several milestone books found their way to the shelves (notably, Anderson, 2014; Bechtel, 2008; Craver, 2007); and the number of papers on neuroscientific topics in philosophy of science journals increased (Malaterre et al., 2019; 2020). Meanwhile, a neuroscience-oriented naturalization affected several fields of philosophy, from epistemology to ethics (Knobe, 2015). In some cases, this was acknowledged (and emphasized) by the prefix "neuro", as in neuroethics (Roskies, 2002) and neuroaesthetics (Zeki, 1999). But often, neuroscientific evidence slowly made its way in more "classic" philosophical papers, even without any fancy "neuro" prefix.

According to Bickle (2019), one explanation for the "fall" (his quotation marks) of the neurophilosophy revolution is that "the field began to attract new participants, who were more interested in digging deeper into neuroscience itself and less concerned about revolutionizing perennial philosophy" (p. 8). Indeed, within the generation of scholars trained in the aftermath of the neurophilosophy (failed?) revolution, it is common to find researchers who regularly publish in both philosophical and neuroscientific venues. Current research at the frontier between philosophy and

neuroscience is in fact characterized by “highly technical research projects, where the specific contributions of the scientists and the philosophers grow increasingly intertwined and difficult to distinguish” (Bickle, 2019, p. 7). Examples of debates in philosophy of neuroscience that have their roots in methodological papers that appeared in neuroscience journals include the philosophical debate on cognitive ontology (Janssen et al., 2017; McCaffrey, 2023), which can be traced back to a paper published in *Cognitive Neuropsychology* (Price & Friston, 2005), and that on reverse inference (e.g. Calzavarini & Cevolani, 2022; Nathan & Del Pinal, 2017), which stem from an article in *Trends in Cognitive Sciences* (Poldrack, 2006).

Several facts may be invoked that either showcase or explain the increased porousness between philosophy and neuroscience (or both). A case at hand is the rapid development of the field of neuroethics. Within only a couple of decades since the first documented use of the term in a 2002 meeting by the Dana Foundation (Roskies, 2002), neuroethics has managed to become a well-established disciplinary sub-field, with its own teachings, scientific societies (e.g. the International Society for Neuroethics), and journals (most notably *AJOB Neuroscience* and *Neuroethics*). Other events worth mentioning are the initiatives started in the 2010s with the explicit aim to incubate these dialogues between philosophy and neuroscience, such as the Summer Seminars in Neuroscience and Philosophy, held annually at the Duke University since 2016; and the Neural Mechanisms Online webinars, starting in 2018. Moreover, new attempts in creating dedicated venues for both philosophers and neuroscientists have been recently revamped, with the foundation of journals such as *Neuroscience of Consciousness* (2015), and *Philosophy and the Mind Sciences* (2020). While it is too early to assess their impact, the very fact that both journals have already outlived the aforementioned *Brain and Mind*, together with the other cues described in this section, suggest that the interaction between the two disciplinary fields has grown significantly stronger.

3 Research questions and hypotheses

The portray sketched above suggests that in the last decades, a fertile “trading zone” (Galison, 1997) has thrived between neuroscience and philosophy. In order to better understand how this interaction has unfolded and how it is currently developing, in the present paper we aim to address the following research questions:

R1: Visibility over time. On the background of a world with increasingly interest in the brain, how has the philosophers’ interest in neuroscience developed over time? Was it matched by a corresponding attention from neuroscientists toward philosophy?

R2: Neuroscience in Philosophy. Academic philosophy includes numerous areas (metaphysics, philosophy of science, ethics, history of philosophy, etc.). Is the visibility of neuroscience homogeneously distributed among the different areas of philosophy or is it concentrated in few, specific provinces (e.g., philosophy of science or philosophy of mind)?

R3: Philosophy in Neuroscience. Neuroscience too encompasses several areas, such as cognitive neuroscience, neurology, molecular and cellular neuroscience, and neuropsychiatry. Is philosophy equally visible across all these areas, or are there specific sub-areas of neuroscience where philosophy is more prominent?

R4: The trading zone. How is the trading zone between neuroscience and philosophy configured? What are the specific points of contact where exchanges between the two fields happen?

Research questions R1-R4 are *explorative* in nature, in the sense that they aim to descriptively map the different dimensions of the philosophy and neuroscience nexus without testing specific claims. However, based both on the historical reconstruction of the previous section and on observations made by experts working at the intersection of philosophy and neuroscience, it is possible to associate to R1-R4 specific *hypotheses*.

H1: A neuro-revolution has occurred in philosophy. If neuroscience had a revolutionary impact on philosophy, as envisioned by neurophilosophers, we should expect a temporal trend characterized by an abrupt increase in the visibility of neuroscience in philosophy (R1) and a diffusion of neuroscience in all main areas of philosophy (R2), including the core areas of (analytic) philosophy (philosophy of language, epistemology, mind and metaphysics, see (Kitcher, 2011)). As a corollary, we could observe some impact of philosophy on neuroscience themselves (R3), as a neuro-contaminated philosophy should hopefully be more relevant for neuroscience.

H2: Psychology is the disciplinary link between philosophy and neuroscience. According to Bickle (2003b), philosophers interested in neuroscience have focused almost entirely on the “higher layers” of brain organization, i.e., the so-called cognitive or systems neuroscience. “Lower layers”, i.e., cellular and molecular neuroscience, in contrast, have been neglected by philosophers. If Bickle’s assessment is correct, we should expect that the trading zone between philosophy and neuroscience (R4) is mainly populated by psychologically-oriented topics, and that the neuroscience which is more visible (R2) and more interested (R3) in philosophy is cognitive neuroscience instead of the biomedical branches.

H1 and H2 serve as useful guides for interpreting the results of the exploratory questions R1-R4. However, we emphasize that answering R1-R4 is intrinsically valuable, as a detailed quantitative mapping of the neuroscience-philosophy nexus is currently lacking. Even if not all the results can be interpreted in light of H1 and H2, they still provide valuable material for reflecting on the state of the connection between the two fields.

4 Methodology

Both R1-R4 and H1-H2 focus essentially on processes of communication and interaction between *entire disciplines*. As such, their scope extends over dozens, if not thousands, of scientific publications and researchers, over several decades.

Traditional methods in philosophy of science, such as case-studies and conceptual analysis, are ill-suited for investigating phenomena of this size (Pence & Ramsey, 2018; Ramsey & Block, 2022). A quantitative approach based on *scientometrics*, by contrast, offers methodologies able to cope with potentially millions of publications (Boyack & Klavans, 2019; Mingers & Leydesdorff, 2015).

In particular, our study will focus on *scientific journals* and *citations*, which constitute the building blocks of our quantitative analyses.

Journals have played a key role in the communication system of modern science since the Nineteenth century (Baldwin, 2015; Csiszar, 2018). Even in the humanities, and notably in philosophy, journals are acquiring an increasing role both as publication venues (Larivière et al., 2006) and as *loci* of gatekeeping and prestige attribution (Katzav & Vaesen, 2017a, 2017b). From a bibliometric perspective, moreover, journals represent a convenient meso-scale in between individual publications and larger classificatory units, such as disciplines or scientific fields (Sugimoto & Weingart, 2015). We think therefore that journals offer the right scale to observe the kind of phenomena we are interested in.

Citations, on the other hand, are the inter-textual links that connect scientific publications via the cited references listed in bibliographies (Hyland, 1999). Citations have been used since the 1960s to reconstruct networks of scientific publications and build quantitative maps of scientific fields (Börner, 2010; Petrovich, 2021; Small, 2003). In this study, we will consider citations as *traces of communication* between research areas. Following Khelifaoui et al., (2021), we prefer to construe these traces via the concept of “visibility” instead of using epistemically committed notions such as “knowledge flow” because we recognize that citations do not always correspond to knowledge transfer (as in the case of the so-called perfunctory citations, see Moravcsik & Murugesan, 1975) and may even be associated with knowledge contestation (negative citations, see Petrovich, 2018). However, we maintain that citations can be used as effective proxies of communication process between research fields when relatively massive bibliometric entities such as journals are considered. Citation flows between journals, in fact, involve usually a substantive number of citations, in which the weight of anomalies such as perfunctory and negative citations turns out to be negligible (Leydesdorff et al., 2011; Van Raan, 1998).

What is more, citation analysis, compared to other quantitative methods, such as topic modelling (e.g., Malaterre et al., 2019), is particularly well suited for investigating the interaction between disciplines as it allows us to track the citation flows between fields and to build citation-based science maps that can illuminate structural relationships and interdependences (Khelifaoui et al., 2021; Petrovich, 2021).

In order to answer R1-R4 and to test H1-H2, we will leverage scientific journals and citation analysis in two ways. First, we are going to use citations to compute *indicators of visibility* (Sect. 4.3 below), which can measure how much a field (e.g., philosophy) is visible in terms of citation impact within another field (e.g., neuroscience). Second, we are going to use citation relationships between journals to build *science maps* of our target fields, philosophy and neuroscience (Sect. 4.4). Such maps will allow us to localize the spots where the visibility of the other field is particularly prominent and to chart the trading zone. Before explaining in details how indicators are defined and science maps built, however, we need to present our

data source (Sect. 4.1) and discuss how we operationally defined philosophy and neuroscience in bibliometric terms (Sect. 4.2).

4.1 Data

All our analyses are based on data extracted from Clarivate Analytics Web of Science (WoS) database (<https://www.webofscience.com/>).¹

We preferred Web of Science (WoS) over its main commercial competitors (Elsevier's Scopus and Digital Science Dimensions) because of its temporal coverage and established usage in prior bibliometric studies of philosophy-science interactions (Bonino et al., 2022; Khelifaoui et al., 2021; McLevey et al., 2018; Pradeu et al., 2021; Yan & Liao, 2023) and in citation studies of philosophy (Chi & Conix, 2022; Noichl, 2019; Petrovich & Buonomo, 2018). Even if Scopus coverage of the humanities and social sciences is higher than WoS', Scopus data are not reliable before 1996 (Sugimoto & Larivière, 2018). Scopus, therefore, cannot support the desired temporal scope of the analysis. Experiments with Dimensions, similarly, show that data for records before the 1990s remain limited or incomplete – though its coverage is rapidly improving.² As to open-access citation databases (e.g., CrossRef or the recently launched OpenAlex), to the best of our knowledge, they have not yet been used for citation analysis of philosophy. Hence, we do not have sufficient data on their reliability and coverage.

4.2 Bibliometric delineation

Web of Science covers both scientific journals and book collections (but not books). Each of them is assigned to one or more of the 254 Web of Science Categories, based on its subject matter and scope, author and editorial board affiliations, citation relationships with other journals, the journal's sponsor (e.g., a scientific society), and journal's classification in other bibliographic databases.³ All the publications appeared in a journal or book collection inherit the categories of the journal.

Journals and book collections publishing research in philosophical topics are assigned to four WoS Categories: 'Philosophy', 'History & Philosophy of Science', 'Ethics', and 'Medical Ethics'. 1517 journals or book collections are classified in these categories, for a total of around 730,000 individual publications. Journals and book collections publishing research in neuroscience, on the other hand, are

¹ Queries were made through the SQL relational database system hosted by the Centre for Science and Technology Studies (CWTS) at Leiden University, using the 2024 version of Web of Science.

² Another shortcoming of Dimensions is that it does not use a journal-level classification system, like Web of Science and Scopus. Instead, it employs a publication-level classification system, where each publication is assigned to one or more disciplinary categories. Dimensions' classification system is based on a combination of existing classificatory schemes and artificial intelligence. In our opinion, at its current stage of development, this system still produces suboptimal assignments for philosophy, resulting in too many false positives and false negatives.

³ For more details, see https://support.clarivate.com/ScientificandAcademicResearch/s/article/Web-of-Science-Core-Collection-Web-of-Science-Categories?language=en_US

assigned to two WoS Categories: ‘Neurosciences’ and ‘Neuroimaging’. 612 journals or book collections are classified in these categories, for a total of around 1,900,000 individual publications.

In the following analyses, we will call PHI_L (“Philosophy large”) the set of 1517 journals or book collections that belong to the four philosophically relevant WoS categories. Similarly, we will call NEU_L (“Neuroscience large”) the set of 612 journals or book collections that belong to the two neuroscientific WoS categories.

Now, not all the journals or book collections in PHI_L and NEU_L are equally covered by Web of Science. For instance, only a limited number of citations that occur in the French philosophy journal *Revue philosophique de la France et de l'étranger* point to publications that are in turn covered by Web of Science, because of the well-known poor coverage of non-English speaking outlets (Sugimoto & Larivière, 2018). The citation analysis that can be performed on the *Revue* in WoS is therefore significantly limited and its results are likely to be skewed. What is more, not all the publications appeared in journals and book collections belonging to PHI_L and NEU_L are scientifically relevant: editorials, letters, discussions, meeting abstracts, biographical items, and book reviews, while technically publications, do not convey the same amount of scientific information as research articles.

In order to focus on the items that are both scientifically relevant and sufficiently covered in WoS, we define therefore two subsets of PHI_L and NEU_L : PHI_N (“Philosophy narrow”) and NEU_N (“Neuroscience narrow”), respectively. They are constructed based on the following criteria: a) they include scientific journals only, no book collection; b) only publications classified as research article, literature review, or conference proceeding (in the following, we will call these publications “research documents”) are included⁴; c) only research documents published after 1980 are included; d) only journals with at least 300 research documents are retained; e) only journals with at least 500 citations that point to other research documents covered in Web of Science are retained; f) only journals with at least 10% of their total citations that point to other research documents in Web of Science are retained. The two last criteria exclude the case of journals, such as the *Revue*, for which the coverage of WoS is so limited that any citation analysis would be methodologically unsound.

Table 1 reports the descriptive statistics of PHI_N and NEU_N . Note the significant difference between the two in terms of the number of citations that point to other publications in WoS. Only 32% of citations appearing in PHI_N point to publications in WoS, against the 87% of NEU_N . One of the main causes of this difference is the considerable weight that books and monographs still have in philosophy. Since books and monographs are not covered by WoS, all citations that refer to them point in fact to materials outside the database. For example, a citation to Kant’s *Critique of pure reason* counts as an out-going citation, but not as a *recorded* out-going citation because it refers to a document that lies outside WoS.

⁴ Document types used in WoS are explained at <https://webofscience.help.clarivate.com/en-us/Content/document-types.html>

Table 1 Descriptive statistics of PHI_N and NEU_N . Research documents include only publications classified as research article, literature review, or conference proceeding. Recorded out-going citations is the number of citations that point to a publication that is covered in Web of Science

Set (# journals)	WoS Categories (# journals)	Publications	Research does (% of total)	Tot out-going citations	Recorded out-going citations (% of total)
PHI_N (231)	Philosophy (135), History & Philosophy of Science (62), Ethics (50), Medical Ethics (16)	331,361	199,517 (60.2%)	7,302,369	2,363,114 (32.4%)
NEU_N (321)	Neuroimaging (18), Neurosciences (310)	1,687,666	1,142,041 (67.7%)	58,332,442	50,822,687 (87.1%)

4.3 Definition of the visibility indicators

As explained above, we consider citations as bibliometric evidence of interactive processes between research fields: if research field x interacts with research field y , we expect that publications in x cite publications in research field y , i.e., that research field y is citation-wise *visible* in research field x .

To measure this visibility, we propose the following *visibility indicator* V :

$$V(x, y)_\alpha = \frac{\sum_{i=1}^{d_\alpha} v_{i\alpha}}{d_\alpha}$$

V is computed for two research fields x and y , where y is the research field of which we are measuring the visibility in x , and a unit of analysis α , which can be any bibliometric entity, intended as a collection of publications. As explained above, in this study, we will set in most of the cases α as a scientific journal, but other α can be easily conceived (e.g., the collection of publications of an author, the set of all publications published in a certain year, etc.). The instances of α , i.e., the different journals in our case, belong to the research field x according to some classificatory system – in our case, the system of WoS Categories presented above. The number of publications belonging to each instance of α is represented by d_α : in our case, d_α is the number of research documents published by a certain journal classified in x . Lastly, $v_{i\alpha}$ is a binary variable that equals 1 if the i -th publication in the target instance of α cites at least one publication in research field y , and 0 otherwise.

In non-technical terms, V is nothing else than the *ratio* of the research documents published in a unit of analysis classified in research field x that cite at least one publication that is classified in the research field y , over the total number of research documents in the considered unit of analysis. This ratio is of course bounded in the range $[0,1]$ and can be expressed as a percentage.

The visibility indicator of neuroscience in philosophy is therefore:

$$V(PHI_N, NEU_L)_j = \frac{\sum_{i=1}^{d_j} v_{ij}}{d_j}$$

where j is a journal belonging to PHI_N as defined above. By the same token, the visibility indicator of philosophy in neuroscience is:

$$V(NEU_N, PHI_L)_k = \frac{\sum_{i=1}^{d_k} v_{ik}}{d_k}$$

where k is a journal belonging to NEU_N as defined above.

Note that in both declinations of V , the research field y is defined as the *extended set* and not as the narrow set of scientific journals. To justify this asymmetry, it is convenient to make an example. We have seen above that the *Revue philosophique de la France et de l'étranger* is excluded from PHI_N because of the insufficient WoS coverage. We will not therefore compute the indicator $V(PHI_N, NEU_L)_j$ when $j =$ *Revue philosophique de la France et de l'étranger* as the score on the indicator is

highly likely to be unreliable. However, if a research document published in the neuroscience journal k cites the *Revue*, we want the binary variable v_{ik} in $V(NEU_N, PHI_L)_k$ to take the value of 1, as that citation still indicates an interaction with philosophy, which we do not want to overlook. This is why we always define the research field whose visibility we are measuring as the extended set instead of the narrow set.

Another key property of V is that the value of the binary variable $v_{i\alpha}$ depends on the *direction* of the citation flow. In fact, the variable takes the value of 1 when the i -th publication from field x *cites* at least one publication from field y . This means that $v_{i\alpha}$ is sensitive essentially to *out-going citations* from x to y . However, it is possible to use an alternative binary variable, $w_{i\alpha}$, that takes the value of 1 when the i -th publication from field x *is cited* by at least one publication from field y . $w_{i\alpha}$ captures hence *incoming citations* from y to x .

We call the resulting indicator W *inverse visibility*:

$$W(x, y)_\alpha = \frac{\sum_{i=1}^{d_\alpha} w_{i\alpha}}{d_\alpha}$$

Of course, there is no reason to assume a priori that V and W are correlated, as the flow of out-going citations from x to y does not necessarily reflect the flow of incoming citations from y to x . For instance, a journal in philosophy could cite a lot of publications in neuroscience (high V) without being cited a lot by neuroscience journals in turn (low W), or vice versa.

Inverse visibility for philosophy therefore measures, for each journal in PHI_N , the percentage of its publications that receive at least one citation from neuroscience (NEU_L):

$$W(PHI_N, NEU_L)_j = \frac{\sum_{i=1}^{d_j} w_{ij}}{d_j}$$

And inverse visibility for neuroscience measures, for each journal in NEU_N , the percentage of its publications that receive at least one citation from philosophy (PHI_L):

$$W(NEU_N, PHI_L)_k = \frac{\sum_{i=1}^{d_k} w_{ik}}{d_k}$$

Of course, alternative indicators of visibility (and inverse visibility) can be conceived. We could measure the ratio between the number of citations in x that point to y over the total number of citations in x . Or, we could stipulate that $v_{i\alpha}$ takes the value of 1 only if the i -th publication in the target instance of α cites the field y above a certain threshold, i.e., it cites at least n publications from the field y (or at least $n\%$ of its citations point to field y). In the Online Supplementary Materials, the indicator based on the ratio of citations is reported, and (Petrovich & Viola, 2024) report results based on the citation threshold method. We note that the overall results remain similar, independently of the particular design of the visibility indicator

chosen. Therefore, in order to simplify the presentation, we are going to focus only on V and W as defined above.

Starting from V and W , more refined indicators can be easily constructed. For instance, if we have a set of journals J (e.g., all journals in a certain research field), we can compute the value of V for J as the *weighted average* of the V of the individual journals belonging to J :

$$\overline{V(x, y)}_J = \sum_{j=1}^{c_J} \frac{d_{jJ}}{d_J} \cdot \frac{\sum_{i=1}^{d_{jJ}} v_{ij}}{d_{jJ}}$$

where c_J is the number of journals in J (i.e., $|J|$), d_J is the total number of research documents published in journals belonging to J , d_{jJ} is the number of research documents in journal $j \in J$, and $\frac{d_{jJ}}{d_J}$ is the proportion of journal j research documents over the total of research documents published in journals belonging to J . The average value of W for J is defined in the same way, only replacing v with w .

The last indicators we are going to define are simple measures of the *incidence* of a research unit α in a certain reference set. For instance, if α is the research field r and the reference set is the whole WoS database, the incidence indicator for r will be defined as:

$$I_r = \frac{d_r}{d_w}$$

where d_r is the total number of research documents published in journals classified in the research field r , and d_w is the total number of research documents in WoS. This indicator is useful as a *baseline* to which we can compare the visibility of research fields. If, for instance, the visibility of neuroscience in philosophy is much higher than the overall incidence of neuroscience in WoS, we can conclude that the visibility of neuroscience in philosophy is higher than expected.

If we set α to a subset g of research documents in r , lastly, we define the *relative incidence* of the subset g as the ratio between d_g (the research documents classified in g) and d_r :

$$I_g = \frac{d_g}{d_r}$$

4.4 Citation-based journal maps

To answer research questions R2-R4, we do need only reliable indicators of visibility and incidence. We need also to reconstruct the *internal organization* of both philosophy and neuroscience. As noted in Sect. 3, in fact, both fields are divided into several sub-disciplines or sub-areas. The main issue, especially for philosophy, is that there is no consensus on what these sub-areas are or where their boundaries lie (see contributions in Allen & Beavers, 2011). Any representation of philosophy

inevitably incorporates assumptions about the nature, methods, and scope of philosophy, which remain open to debate (Rescher, 2005).

Without denying the importance and philosophical interest of these discussions, in this study we approach the problem of *mapping* philosophy (and neuroscience) empirically, leveraging the techniques of citation mapping, which have been proven to be rather effective in previous attempts of charting the field (Higgins & Smith, 2013; Noichl, 2019; Petrovich, 2024; Petrovich & Buonomo, 2018).

Specifically, we will use journals as units of our maps and measure the similarity between journals based on the analysis of their citation profiles, i.e., the citations that are made by the research documents they publish. In particular, we use the degree of *overlap* between the out-going citations of two journals as measure of their similarity. This technique, known as *journal bibliographic coupling* (Kessler, 1963; Petrovich, 2021), is based on the idea that, if the articles published in two journals have many cited references (i.e., out-going citations) in common, then it is likely that these two journals are intellectually close, e.g., that they publish research on similar topics. By contrast, if the articles of our target journals have no cited reference in common, then they are likely to belong to different research fields. Measuring in this way the similarities between each pair of journal, we can build a *network*, in which nodes represent journals and the weight of the links between journals measures the bibliographic coupling strength between them (Waltman & van Eck, 2014).⁵ In such a network, similar journals will tend to cluster together, forming denser areas that correspond to intellectual units, such as sub-fields or specialized areas. For instance, we expect that in the philosophy journal map, philosophy of science journals will form a group distinct from the cluster of ethics journals. Similarly, journals in cognitive neuroscience and journals in molecular neuroscience should be placed in different clusters of the neuroscience journal map. A clustering algorithm can then be used to automatically detect such groups of similar journals.

The most important feature of journal bibliographic coupling, for our aims, is that it allows us to reconstruct the internal organization of philosophy and neuroscience without relying on external classifications, which are to a certain extent inevitably subjective, but considering only the *objective citation profiles* of journals. The similarity between journals in this approach results from the aggregated citation behavior of all the authors that contribute to those journals, who collectively determine the citation similarity between journals (Kreuzman, 2001). It is not based on a priori taxonomies.

To produce our journal maps of philosophy and neuroscience, we used the science mapping software VOSviewer (van Eck & Waltman, 2010). VOSviewer takes as input the citation profiles of the target bibliometric units (in our case, the journals) and produces as output network visualizations in which the distance between the nodes (i.e., the journals) reflects their similarity, so that similar journals will be

⁵ More precisely, we measured the similarity between journals using the *cosine similarity* between their vectors in the out-going citations space. Cosine similarity has the significant advantage, compared to raw co-occurrence frequencies, of allowing comparisons between journals of different sizes (Salton & McGill, 1983; van Eck & Waltman, 2009).

placed closer and dissimilar journals far apart (van Eck et al., 2010). In this way, the spatial organization of the VOSviewer map reflects the epistemic organization of the field.

VOSviewer is also equipped with a clustering algorithm that automatically attributes journals to different clusters based again on their citation similarity (Waltman et al., 2010). As explained above, these clusters correspond to groups of journals that are characterized by similar citation profiles.⁶

4.5 Summary scheme of the methodology

Now that we have presented in details all the components of our methodology, it is convenient to show how they allow us to address the research questions R1-R4 and test the hypotheses H1-H2.

To answer R1, we are going to compute the average values of $V(PHI_N, NEU_L)_j$ and $V(NEU_N, PHI_L)_k$ for each year in the observation window available for our access to WoS data (1980–2023).⁷ The resulting time-series will allow us to track the changing visibility of neuroscience in philosophy and of philosophy in neuroscience. To have a benchmark, we are going to compare these trends with the changing incidence of philosophy and neuroscience in the overall WoS measured via the incidence indicator.

To answer R2, we are going to combine the map of philosophy journals generated with bibliographic coupling with the visibility indicator $V(PHI_N, NEU_L)_j$, computed for each journal on the map. The visualization will allow us to individuate the journals with the highest value of visibility and locate them in the overall structure of the field. Computing average values of the indicator by cluster will allow us to understand in which area of philosophy neuroscience is more visible.⁸

By the same token, to answer R3 we are going to combine the map of neuroscience journals with the visibility indicator $V(NEU_N, PHI_L)_k$, computed for each journal on the map. Again, the visualization will allow us to identify neuroscience journals where philosophy is particularly visible and to understand in which area of neuroscience the impact of philosophy is higher.

⁶ Note that the algorithm is governed by a *resolution parameter*: the higher the parameter, the more clusters will tend to be individuated, the smaller the parameter, the less clusters will tend to be individuated. There is no golden truth for the choice of the resolution parameter, which should be adjusted by the analyst based on their knowledge of the investigated domain (Petrovich, 2021).

⁷ As explained above, we are going to use a weighted average that takes into account the different size of journals. Formally, the value of the visibility indicator in the year y is computed as:

$$\overline{V(PHI_N, NEU_L)_y} = \sum_{j=1}^{c_y} \frac{d_{jy}}{d_y} \cdot \frac{\sum_{i=1}^{d_{jy}} v_{ij}}{d_{jy}}$$

where c_y is the number of journals in PHI_N that are active in year y and $\frac{d_{jy}}{d_y}$ is the proportion of journal j research documents published in the year y over the total number of research documents in year y .

⁸ As above, we are going to use a weighted average that takes into account the different weight of journals in the cluster to compute average values for clusters.

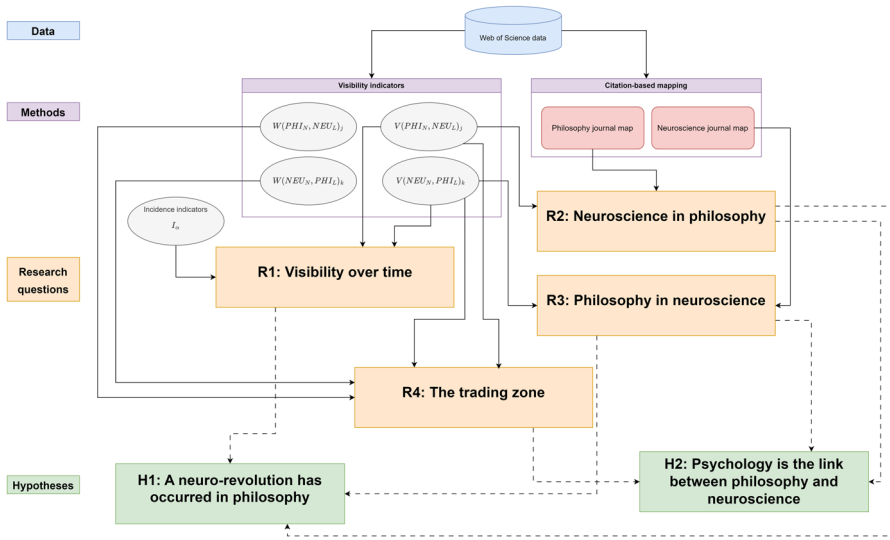


Fig. 1 Diagram summarizing the relationships between data, methodology, research questions, and hypotheses

Lastly, to answer R4 we are going to combine the indicators of visibility and inverse visibility to describe the dynamics of the trading zone between philosophy and neuroscience. In particular, we assume that journals that are characterized by high values of both V and W function as *active loci* of bidirectional communication. As crossroads of both incoming and outgoing citations, they act as *brokers* of the citation flow between philosophy and neuroscience.

As to the hypotheses, combining results of R1-R3 (temporal trends together with the structural information obtained from the maps) will allow us to test H1. To evaluate H2, on the other hand, we are going to combine the results of R2-R4 together with a specific analysis of the inverse visibility of neuroscience in philosophy. If cognitive neuroscience is disproportionately more visible than molecular-cellular neuroscience in philosophy, i.e., if journals most cited in philosophy belong to cognitive neuroscience and they are cited more often than those belonging to molecular-cellular neuroscience (normalized based on their incidence in neuroscience), then the hypothesis is verified.

The diagram in Fig. 1 summarizes how the different components of the methodology (visibility indicators, incidence indicators, and citation maps) relate to our research aims (R1-R4 and H1-H2).

5 Results

The left panels of Fig. 2 show the average visibility of neuroscience in philosophy over time ($\overline{V(PHI_N, NEU_L)_{PHI}}$, Fig. 2A), compared with the incidence of neuroscience in WoS (I_{NEU} , Fig. 2C) over the same time window. The two trends are not

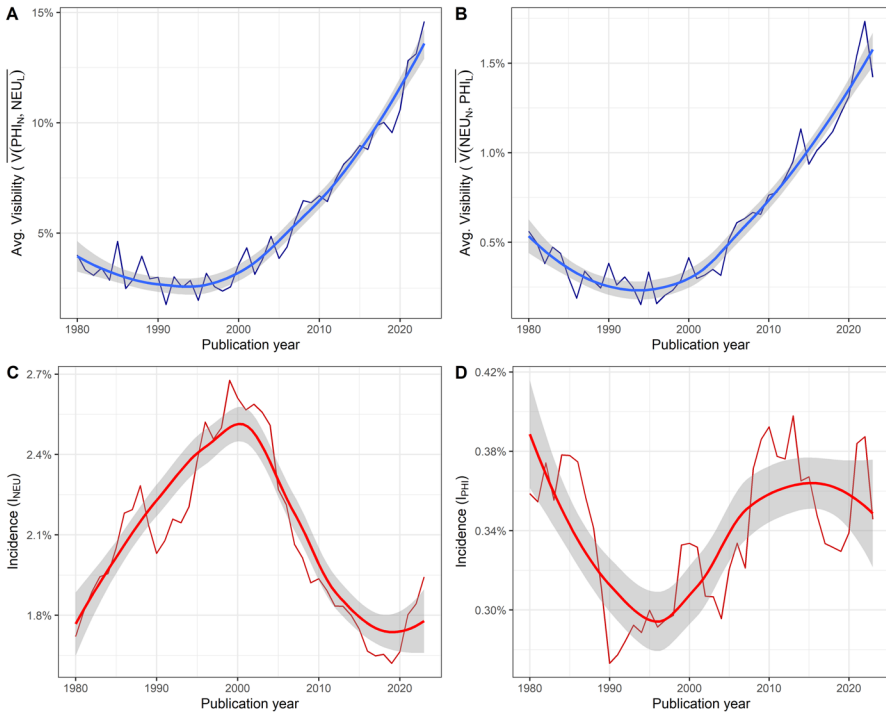


Fig. 2 Visibility and incidence over time. Panel A: Average visibility of neuroscience in philosophy; Panel B: Average visibility of philosophy in neuroscience; Panel C: Incidence of neuroscience in WoS; Panel D: Incidence of philosophy in WoS. Note the different scales of the y-axes. Average visibility is computed as the weighted average of the visibility of all journals in a certain field for the year y , see Footnote 7

correlated. In fact, the first two decades of the observation period show an increase in the incidence of neuroscience in the database from 1.8% to 2.7%, which is not matched by a parallel increase of the visibility of neuroscience in philosophy, that remains stable at around 3%. After the turn of the century, neuroscience's incidence in WoS decreases, whereas the visibility of neuroscience in philosophy increases linearly, reaching a peak of 14.6% in 2023. The increasing attention of philosophers for neuroscience starting from 2000, therefore, cannot be attributed to an increase of the weight of neuroscience in the database, as the latter in fact has decreased.

The right panels of Fig. 2 show the average visibility of philosophy in neuroscience over time ($v(PHI_N, NEU_L)_{NEU}$, Fig. 2B), compared with the incidence of philosophy in WoS (I_{PHI} , Fig. 2D) in the same period. Again, the two trends are not correlated. The incidence of philosophy in the database has remained overall stable, oscillating around an average of 0.3%. The visibility of philosophy in neuroscience, on the other hand, has started to increase significantly at the turn of the century, raising from 0.4% to 1.5% in 2023. As before, such an increase cannot be explained by a change in the weight of philosophy in the database. Interestingly,

the two visibility trends (Fig. 2A-B) are correlated, as both begin to increase after 2000. Note, however, that they differ of an order of magnitude.

Figure 3 shows the journal map of philosophy. Interestingly, the spatial organization of journals reflects to a certain extent a division between theoretical vs applied philosophy: journals with a more theoretical orientation are on the south-western side of the map, whereas applied journals are concentrated in the north-eastern corner. At a finer-grained level, the algorithm detects 9 clusters of journals, which correspond more or less to philosophical specialties. We overlaid labels to clusters on the map to ease interpretation but remember that these labels should be taken with a grain of salt. These labels – that we will express in small caps – are a convenient way to refer to clusters but should not be intended as definitive classificatory categories for philosophy journals or philosophy sub-disciplines (see Petrovich, 2024, Chapter 4 for a detailed discussion of labelling in bibliometric clustering). Starting from the north-eastern corner, we find the green cluster including journals in different branches of APPLIED ETHICS (*Medicine Health Care and Philosophy, Journal of Business Ethics, Science and Engineering Ethics*). Then, moving counterclockwise, we meet the purple cluster of journals in MORAL AND SOCIAL PHILOSOPHY (*Ethics, Ethical Theory and Moral Practice, Law and Philosophy, Philosophy & Public Affairs*), and two small clusters, an orange one including journals in PHILOSOPHY OF RELIGION AND ORIENTAL PHILOSOPHY (*American Catholic Philosophical Quarterly, Asian Philosophy*), and a violet one with journals in AESTHETICS and philosophy of arts (*British Journal of Aesthetics, Journal of Aesthetics and Arts Criticism*). Then, moving to the western side of the map, we meet a big blue cluster comprising all mainstream journals in ANALYTIC PHILOSOPHY (*Mind, Nous, Philosophical Review, Journal of Philosophy, Philosophy and Phenomenological Research*). These journals publish the

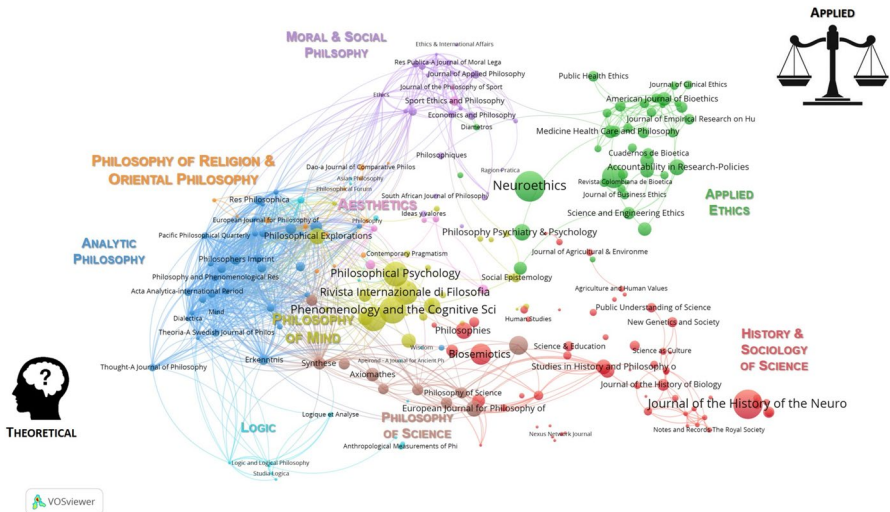


Fig. 3 Journal map of philosophy. Each node represents a journal in the PHI_N set ($n = 231$). The size of the nodes is proportional to the value of $V(PHI_N, NEU_L)$. The color of the nodes corresponds to the cluster (resolution parameter = 1.3, minimum cluster size = 3). An interactive visualization with further statistics is available at <https://tinyurl.com/2dbpkv8s>

core areas of analytic philosophy mentioned by Kitcher (philosophy of language, epistemology and metaphysics). A peninsula of journals specializing in LOGIC (*Journal of Philosophical Logic*, *Studia Logica*) extends from the continent of analytic journals (light blue cluster). Then, moving to the center, we meet a yellow cluster that includes, interestingly, both journals in PHILOSOPHY OF MIND and mind sciences (*Phenomenology and the Cognitive Sciences*, *Journal of Consciousness Studies*, *Rivista Internazionale di Filosofia e Psicologia*) and journals on pragmatism (*Transactions of the Charles S Pierce Society*). Lastly, in the south, the brown cluster comprising journals in PHILOSOPHY OF SCIENCE (*Synthese*, *British Journal for the Philosophy of Science*, *Philosophy of Science*) forms a southern corridor that ends in the red cluster of HISTORY AND SOCIOLOGY OF SCIENCE journals (*Social Studies of Science*, *Isis*, *Studies in History and Philosophy of Science*).

Much more could be said on this map, but, in line with the focus of the present paper, we will limit to highlight that neuroscience is not equally visible in all clusters of the map (remember that the size of the dots representing journals is proportional to their score on the $V(PHI_N, NEU_L)$ indicator). As Table 2 shows analytically, the cluster where neuroscience is more visible is the yellow cluster (average visibility = 16.2%) where journals with high neuroscience visibility are concentrated (*Phenomenology and the Cognitive Sciences*, 61.1%, *Journal of Consciousness Studies*, 52.4%, *Philosophical Psychology*, 50.8%). The second cluster with the average highest visibility is the brown cluster of philosophy of science (average visibility = 13.6%). The other clusters show a visibility lower than 10%. Nonetheless, some of them include specific spots of neuroscience visibility, such as the *Journal of the History of the Neurosciences* (71.4%, in the red cluster), and *Neuroethics* (75.7%, in the green cluster, which is, unsurprisingly, the journal with the highest neuroscience visibility in the entire map).

Looking at the inverse visibility, i.e., the visibility that the philosophy journals have in neuroscience, the brown cluster ranks first, with an average inverse visibility of 8.3%, followed by the yellow cluster and the green cluster (both with an average inverse visibility of 5.7%). Again, some journals stand out with significantly higher inverse visibility scores: *Neuroethics* (36.9%), *Journal of Consciousness Studies* (25.5%), *Phenomenology and the Cognitive Sciences* (20.6%).

Figure 4 shows the journal map of neuroscience. The overall configuration of the map is organized around two axes. Moving along the horizontal axis from west to east, the focus of the journals shifts from basic to clinical research. Moving along the vertical axis from south to north, the scale of the investigated biological structures changes from the whole brain to individual neurons – and in some cases even smaller entities at the molecular level, such as neurotransmitters. At a finer-grained level, the clustering algorithm identifies 5 clusters of journals, which can be mapped to different neuroscientific specialties. The green cluster in the upper part of the map includes journals focusing on MOLECULAR AND CELLULAR NEUROSCIENCE (*Synapse*, *Journal of Neurocytology*). Moving to the down and left corner of the map, we find a red cluster of journals of COGNITIVE AND SYSTEMS NEUROSCIENCE (*Trends in Cognitive Science*, *Neuroimage*), including computational modeling (*Neural Computation*). In the central-bottom part of the map, the violet cluster is crowded by journals on NEUROPSYCHIATRY (*Molecular Psychiatry*, *Journal of Psychiatry & Neuroscience*). The blue cluster at the right corner of the triangle is populated by journals dealing with

Table 2 Cluster-level statistics for the philosophy journal map. Formally, research documents (4th column) is $\sum_{j=1}^C d_j$, where C is the number of journals in the cluster (3rd column); research documents citing neuroscience (5th column) is $\sum_{j=1}^C \sum_{i=1}^{d_j} v_{ij}$; average visibility (6th column) is the weighted average of the visibility of the journals in the cluster (see Footnote 8 above); in the 7th column, the three journals with the highest visibility in the cluster are reported, with the number of research documents (d_j) and the value of the visibility indicator of each journal in parentheses; research documents cited by neuroscience (8th column) is $\sum_{j=1}^C \sum_{i=1}^{d_j} w_{ij}$; average inverse visibility (9th column) is the weighted average of the inverse visibility of the journals in the cluster, which is computed in the same way as average visibility; in the 10th column, the three journals with the highest inverse visibility in the cluster are reported, with the number of research documents (d_j) and the value of the inverse visibility indicator of each journal in parentheses

Cluster	Label	Journals	Res. docs	Res. Docs Citing Neurosci.	Avg. Visibility $v(V(PHI_N, NEU_L))$	Top 3 journals with highest V (res. docs, V)	Res. Docs Cited by Neurosci	Avg. Inverse Visibility $W(PHI_N, NEU_L)$	Top 3 journals with highest W (res. docs, W)
1	HISTORY & SOCIOLOGY OF SCIENCE	60	41,980	2,393	5.7%	<i>Journal of the History of the Neurosciences</i> (405, 71.4%) <i>Biosemiotics</i> (392, 37.8%) <i>Philosophies</i> (434, 23.3%)	1,569	3.7%	<i>Journal of the History of the Neurosciences</i> (405, 51.4%) <i>Journal of the History of Medicine and Allied Sciences</i> (687, 8.9%) <i>Social Studies of Science</i> (1273, 8.2%)
2	APPLIED ETHICS	40	46,597	4,555	9.8%	<i>Neuroethics</i> (469, 75.7%) <i>Estetika</i> (767, 42.8%) <i>Philosophy Psychiatry & Psychology</i> (438, 22.1%)	2,645	5.7%	<i>Neuroethics</i> (469, 36.9%) <i>American Journal of Bioethics</i> (594, 16.8%) <i>Perspectives in Biology and Medicine</i> (1,948, 12.0%)
3	ANALYTIC PHILOSOPHY	33	38,674	1,799	4.7%	<i>Philosophy Compass</i> (1,239, 11.9%) <i>Journal of the American Philosophical Association</i> (315, 9.8%) <i>Philosophers Imprint</i> (356, 9.6%)	1,011	2.6%	<i>Journal of Philosophy</i> (1,297, 9.2%) <i>Philosophical Review</i> (541, 7.6%) <i>Philosophy Compass</i> (1,239, 6.1%)

Table 2 (continued)

Cluster	Label	Journals	Res. docs	Res. Docs Citing Neurosci.	Avg. VisibilityV (PHI_N, NEU_L)	Top 3 journals with highestV (res. docs,V)	Res. Docs Cited by Neurosci	Avg. Inverse VisibilityW (PHI_N, NEU_L)	Top 3 journals with highestW (res. docs,W)
4	PHILOSOPHY OF MIND	27	16,881	2,727	16.2%	<i>Phenomenology and the Cognitive Sciences</i> (714, 61.1%) <i>Journal of Consciousness Studies</i> (1,432, 52.4%) <i>Philosophical Psychology</i> (1,318, 50.8%)	961	5.7%	<i>Journal of Consciousness Studies</i> (1,432, 25.5%) <i>Phenomenology and the Cognitive Sciences</i> (714, 20.6%) <i>Philosophical Psychology</i> (1,318, 16.3%)
5	MORAL & SOCIAL PHILOSOPHY	22	15,282	401	2.6%	<i>Ethical Theory and Moral Practice</i> (880, 6.3%) <i>Journal of Applied Philosophy</i> (616, 5.8%) <i>Economics and Philosophy</i> (554, 5.8%)	198	1.3%	<i>Economics and Philosophy</i> (554, 4.7%) <i>Philosophy & Public Affairs</i> (567, 4.1%) <i>Ethics</i> (1,122, 3.8%)
6	LOGIC	18	9,555	119	1.2%	<i>European Journal of Philosophy</i> (950, 4.3%) <i>Anthropological measurements of philosophical research</i> (311, 3.2%) <i>Topicos</i> (367, 2.7%)	65	0.7%	<i>Journal of Philosophical Logic</i> (1,231, 2.2%) <i>European Journal of Philosophy</i> (950, 1.6%) <i>Logique et Analyse</i> (362, 1.1%)

Table 2 (continued)

Cluster	Label	Journals	Res. docs	Res. Docs Citing Neurosci.	Avg. Visibility $V(PHI_N, NEU_L)$	Top 3 journals with highest V (res. docs, V)	Res. Docs Cited by Neurosci	Avg. Inverse Visibility $W(PHI_N, NEU_L)$	Top 3 journals with highest W (res. docs, W)
7	PHILOSOPHY OF RELIGION & ORIENTAL PHILOSOPHY	13	9,937	163	1.6%	<i>European Journal for Philosophy of Religion</i> (492, 3.7%) <i>Dao-a Journal of Comparative Philosophy</i> (391, 3.6%) <i>Filosofskii Zhurnal</i> (493, 2.6%)	56	0.6%	<i>Asian Philosophy</i> (521, 1.3%) <i>Philosophical Investigations</i> (686, 1.3%) <i>Philosophy</i> (1,406, 1.1%)
8	PHILOSOPHY OF SCIENCE	10	14,433	1,960	13.6%	<i>Biology & Philosophy</i> (1229, 28.6%) <i>European Journal for Philosophy of Science</i> (463, 20.7%) <i>Axiomathes</i> (593, 15.0%)	1,201	8.3%	<i>Biology & Philosophy</i> (1229, 17.8%) <i>British Journal for the Philosophy of Science</i> (1230, 14.5%) <i>Philosophy of Science</i> (2,440, 12.5%)
9	AESTHETICS	8	6,178	255	4.1%	<i>Aisthesis</i> (421, 7.6%) <i>Sport Ethics and Philosophy</i> (477, 7.5%) <i>Rivista di estetica</i> (546, 7.1%)	148	2.4%	<i>Journal of Aesthetics and Art Criticism</i> (1,319, 4.1%) <i>British Journal of Aesthetics</i> (1,222, 3.4%) <i>Monist</i> (1,372, 2.6%)

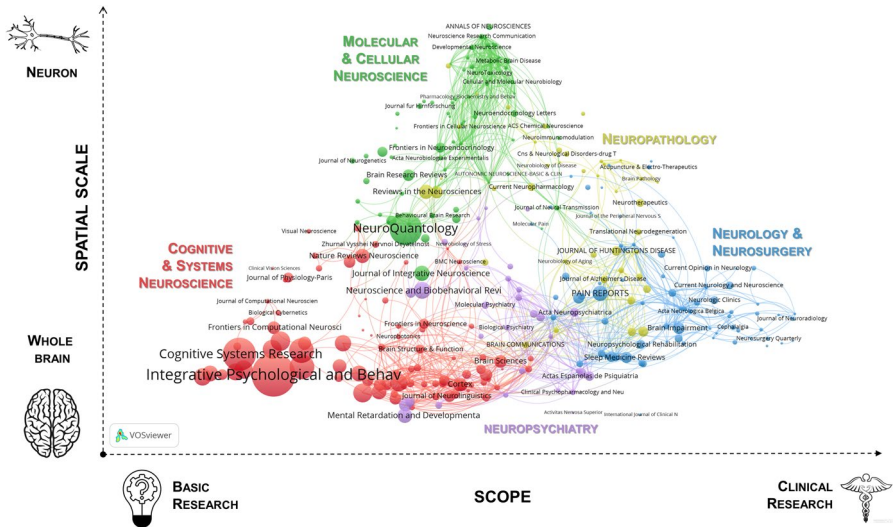


Fig. 4 Journal map of neuroscience. Each node represents a journal. Note that the map shows only the largest component of the journal bibliographic coupling network ($n = 309$), which is a subset of the NEU_N set ($n = 321$). The size of the nodes is proportional to the value of $V(NEU_N, PHI_L)$. The color of the nodes corresponds to the cluster (resolution parameter = 1, minimum cluster size = 3). An interactive visualization with further statistics is available at <https://tinyurl.com/24e7vo5j>

clinical NEUROLOGY AND NEUROSURGERY (*Annals of Neurology, Brain Injury*). Lastly, a smaller yellow cluster, intertwined with the green and the blue one, includes many journals on similar topics, with a shift toward some specific NEUROPATHOLOGY (*Journal of Alzheimer Disease, Neurobiology of Disease*).

Again, we see that philosophy is not homogeneously visible across all these clusters (remember that the size of the dots representing journals is proportional to their score on the $V(NEU_N, PHI_L)$ indicator). As Table 3 shows analytically, the only clusters where philosophy has an average visibility higher than 1% are the red cluster (1.8%) and the violet cluster (1.1%). The journals with the highest philosophy visibility are mainly concentrated in the red cluster (*Integrative Psychological and Behavioral Science*, 37.6%, *Cognitive Systems Research*, 21.9%, *Trends in Cognitive Sciences*, 12.4%) and the journal with the third highest score (*NeuroQuantology*, 21.2%) in in between the red and the green cluster.

Looking at the inverse visibility, i.e., the visibility that the neuroscience journals have in philosophy, we see that the red and violet clusters are again those with the highest average inverse visibility (5.7% and 3.4% respectively).

Moving now to the results concerning the trading zone, Fig. 5 shows two scatter plots. Journals are represented by dots, whose x -coordinate and y -coordinate represent the journal's visibility score and journal's inverse visibility score,

Table 3 Cluster-level statistics for the neuroscience journal map. Formally, research documents (4th column) is $\sum_{j=1}^C d_j$, where C is the number of journals in the cluster (3rd column); research documents citing philosophy (5th column) is $\sum_{j=1}^C \sum_{i=1}^{d_j} v_{ij}$; average visibility (6th column) is the weighted average of the visibility of the journals in the cluster (see Footnote 8 above); in the 7th column, the three journals with the highest visibility in the cluster are reported, with the number of research documents (d_j) and the value of the visibility indicator of each journal in parentheses; research documents cited by philosophy (8th column) is $\sum_{j=1}^C \sum_{i=1}^{d_j} w_{ij}$; average inverse visibility (9th column) is the weighted average of the inverse visibility of the journals in the cluster, which is computed in the same way as average visibility; in the 10th column, the three journals with the highest inverse visibility in the cluster are reported, with the number of research documents (d_j) and the value of the inverse visibility indicator of each journal in parentheses

Cluster	Label	Journals	Res. docs	Res. Docs Citing Phil	Avg. Visibility $V(NEU_N, PHI_L)$	Top 3 journals with highest V (res. docs, V)	Res. Docs Cited by Phil	Avg. Inverse Visibility $W(NEU_N, PHI_L)$	Top 3 journals with highest W (res. docs, W)
1	COGNITIVE & SYSTEMS NEUROSCIENCE	90	286,898	5,239	1.8%	<ul style="list-style-type: none"> <i>Integrative Psycholocial and Behavioral Science</i> (715, 37.6%) <i>Cognitive Systems Research</i> (825, 21.9%) <i>Trends in Cognitive Sciences</i> (1,719, 12.4%) 	16,338	5.7%	<ul style="list-style-type: none"> <i>Trends in Cognitive Sciences</i> (1,719, 55.0%) <i>Nature Reviews Neuroscience</i> (1,418, 35.1%) <i>Nature Human Behaviour</i> (740, 17.7%)
2	MOLECULAR & CELLULAR NEUROSCIENCE	73	419,904	1,641	0.4%	<ul style="list-style-type: none"> <i>NeuroQuantology</i> (793, 21.2%) <i>Journal of Integrative Neuroscience</i> (739, 4.6%) <i>Brain Research Reviews</i> (1280, 2.7%) 	6,861	1.6%	<ul style="list-style-type: none"> <i>Brain Research Reviews</i> (1,280, 10.9%) <i>Neuron</i> (9,979, 9.5%) <i>NeuroQuantology</i> (793, 7.7%)
3	NEUROLOGY & NEUROSURGERY	73	241,846	1,335	0.6%	<ul style="list-style-type: none"> <i>Pain Reports</i> (434, 4.6%) <i>Brain Impairment</i> (373, 2.7%) <i>Journal of Pain</i> (2,609, 2.5%) 	3,837	1.6%	<ul style="list-style-type: none"> <i>Neuropsychology Review</i> (643, 10.6%) <i>Brain</i> (8,005, 9.0%) <i>Pain</i> (9,310, 6.1%)

Table 3 (continued)

Cluster	Label	Journals	Res. docs	Res. Docs Citing Phil	Avg. Visibility (NEU_N, PHI_L)	Top 3 journals with highest V (res. docs, V)	Res. Docs Cited by Phil	Avg. Inverse Visibility $W(NEU_N, PHI_L)$	Top 3 journals with highest W (res. docs, W)
4	NEUROPATHOLOGY	41	94,744	562	0.6%	<p><i>Reviews in the Neurosciences</i> (1,046, 4.0%)</p> <p><i>Progress in Neurobiology</i> (2,225, 2.9%)</p> <p><i>Alzheimer's & Dementia -Translational Research & Clinical Interventions</i> (444, 2.0%)</p>	843	0.9%	<p><i>Progress in Neurobiology Reviews</i> (2,225, 7.1%)</p> <p><i>Neurosciences in the Neurosciences</i> (1,046, 3.9%)</p> <p><i>Alzheimer's Research & Therapy</i> (1,264, 2.5%)</p>
5	NEUROPSYCHIATRY	31	92,234	1,013	1.1%	<p><i>Neuroscience and Biobehavioral Reviews</i> (5,203, 5.9%)</p> <p><i>Mental Retardation and Developmental Disabilities Research Reviews</i> (420, 4.5%)</p> <p><i>Journal of Neurodevelopmental Disorders</i> (547, 3.5%)</p>	3,096	3.4%	<p><i>Mental Retardation and Developmental Disabilities Research Reviews</i> (420, 12.4%)</p> <p><i>Neuroscience and Biobehavioral Reviews</i> (5,203, 11.0%)</p> <p><i>Biological Psychiatry</i> (9,188, 6.3%)</p>

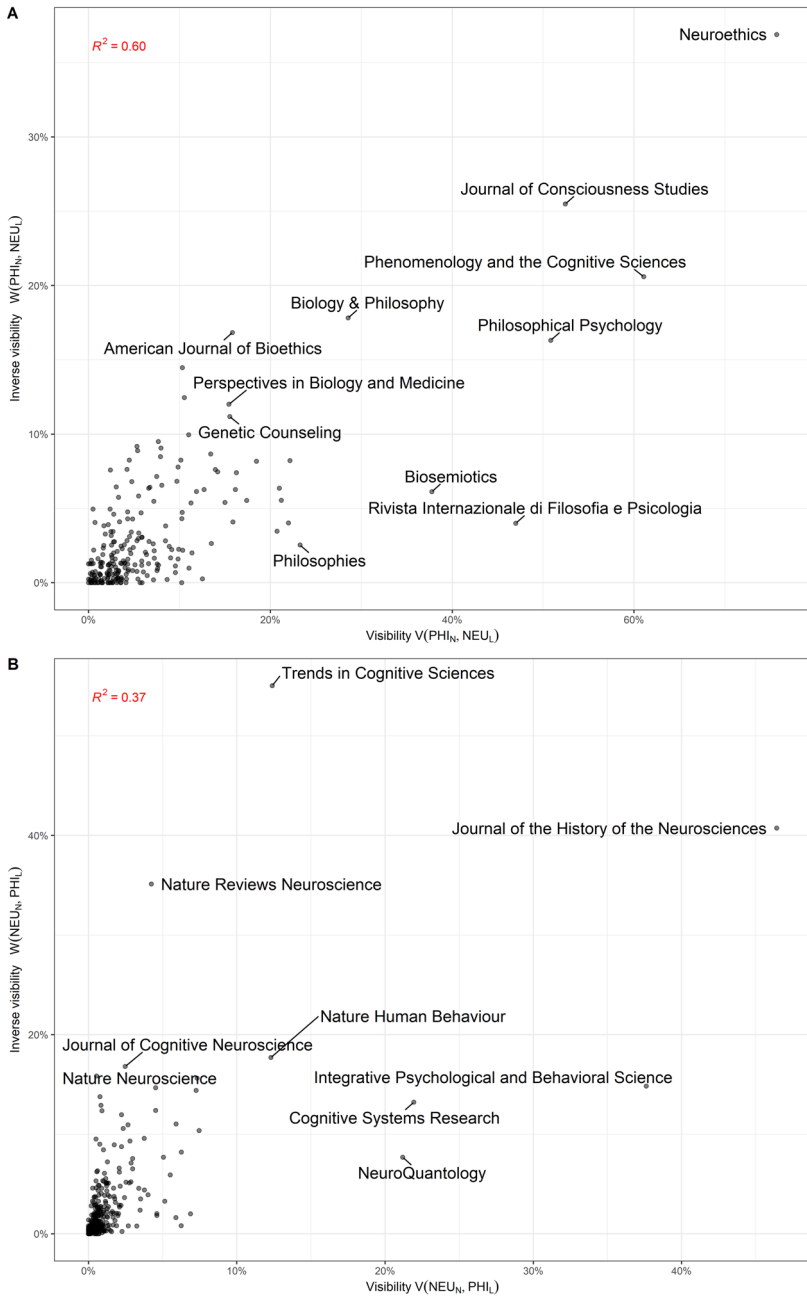


Fig. 5 Journals in the trading zone between philosophy and neuroscience. Panel A: Each dot is a philosophy journals that appears on the map of philosophy (Fig. 3). The x-coordinate is the value of $V(PHI_N, NEU_L)$, the y-coordinate is the value of $W(PHI_N, NEU_L)$. Panel B: Each dot is a neuroscience journal that appears on the map of neuroscience (Fig. 4). The x-coordinate is the value of $V(NEU_N, PHI_L)$, the y-coordinate is the value of $W(NEU_N, PHI_L)$. Labels are shown only for journals with high values of V and W

respectively. Figure 5A represents philosophy journals, Fig. 5B neuroscience journals. The value of the coefficient of determination R^2 is reported in red. In the case of philosophy journals, visibility and inverse visibility show a moderate positive correlation: journals where neuroscience is more visible (e.g., *Neuroethics*, *Journal of Consciousness Studies*) also tend to be visible in neuroscience. In the case of neuroscience journals, by contrast, the weak correlation shows that journals where the visibility of philosophy is relatively higher (e.g., *NeuroQuantology*) are not highly visible in philosophy, and the other way around: journals that have an high inverse visibility, i.e., that are highly cited in philosophy (e.g., *Trends in Cognitive Sciences*, inverse visibility of 55%) are not characterized, in themselves, by a high visibility of philosophy (the philosophy visibility of *Trends in Cognitive Sciences* stops at 12.4%).

Lastly, Table 4 and Table 5 show the statistics of inverse visibility of neuroscience in philosophy.

Table 4 shows the top 15 neuroscience journals with the highest $W(NEU_N, PHI_L)$ and their cluster on the neuroscience map. 13 out of 15 belong to the red cluster, while only 2 belong to the violet cluster. None of the other clusters are represented.

Table 5, on the other hand, compares, for each neuroscience cluster, its *incidence* in neuroscience, measured as the proportion of research documents published in the journals of the cluster over the total number of neuroscience research documents,⁹ with its *relative philosophy visibility*, i.e., the proportion of neuroscience research documents from the target cluster that are cited in philosophy over the total number of neuroscience research documents cited in philosophy.¹⁰ For instance, the journals in the red cluster COGNITIVE & SYSTEMS NEUROSCIENCE publish 25.3% of all research documents published by neuroscience journals. If we consider the neuroscience research documents that are cited in philosophy only, however, the research documents that are published in journals of the red clusters represents 52.7% of them. The difference, in percentage points, between the two is 27.5 p.p., showing that this cluster is over-represented in philosophy, to the extent that its visibility is 27.5 p.p. higher than expected based on the cluster's incidence in philosophy. The other over-represented cluster is NEUROPSYCHIATRY (+1.9 p.p.). All the other clusters, on the other hand, show a relative philosophy visibility which is lower than expected based on their incidence, with the green cluster MOLECULAR & CELLULAR NEUROSCIENCE having the highest negative difference with the baseline.

⁹ Formally, $\frac{\sum_{j=1}^C d_j}{\sum_{C=1}^5 \sum_{j=1}^C d_j}$ where C is the number of journals in the target neuroscience cluster.

¹⁰ Formally, $\frac{\sum_{j=1}^C \sum_{i=1}^d w_{ij}}{\sum_{C=1}^5 \sum_{j=1}^C \sum_{i=1}^d w_{ij}}$, where C is again the number of journals in the target cluster.

Table 4 Neuroscience journals with the highest inverse visibility in philosophy (most cited neuroscience journals in philosophy). Formally, the 5th column is the numerator of $W(NEU_N, PHI_L)$, i.e., $\sum_{i=1}^k w_{ik}$

Rank	Journal	Cluster	$W(NEU_N, PHI_L)$ (%)	Res. docs cited by phil.
1	<i>Trends in Cognitive Sciences</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	55.0	946
2	<i>Nature Reviews Neuroscience</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	35.1	498
3	<i>Nature Human Behaviour</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	17.7	131
4	<i>Journal of Cognitive Neuroscience</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	16.8	686
5	<i>Nature Neuroscience</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	15.8	767
6	<i>Integrative Psychological and Behavioral Science</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	14.8	106
7	<i>Social Cognitive and Affective Neuroscience</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	14.7	285
8	<i>Social Neuroscience</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	14.4	119
9	<i>Cognitive Brain Research</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	13.8	144
10	<i>Cognitive Systems Research</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	13.2	109
11	<i>Current Opinion in Neurobiology</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	12.9	417
12	<i>Mental Retardation and Developmental Disabilities Research Reviews</i>	NEUROPSYCHIATRY	12.4	52
13	<i>Trends in Neurosciences</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	12.4	425
14	<i>Neuropsychologia</i>	COGNITIVE & SYSTEMS NEUROSCIENCE	12.0	1042
15	<i>Neuroscience and Biobehavioral Reviews</i>	NEUROPSYCHIATRY	11.0	574

Table 5 Neuroscience clusters: relative weight in neuroscience vs relative visibility in philosophy. Formally, the 2nd column is computed as the proportion of research documents from the target cluster over the total of neuroscience research documents, i.e., as $\frac{\sum_{j=1}^C d_j}{\sum_{c=1}^S \sum_{j=1}^C d_j}$ where C is the number of journals in the target neuroscience cluster. Clearly, the sum of the column is 100%. The 3rd column is computed as the proportion of neuroscience research documents cited in philosophy from the target cluster over the total number of neuroscience research documents cited in philosophy, i.e., as $\frac{\sum_{j=1}^C \sum_{i=1}^{d_j} w_{ij}}{\sum_{c=1}^S \sum_{j=1}^C \sum_{i=1}^{d_j} w_{ij}}$, where C is the number of journals in the target cluster. Clearly, the sum of the column is 100%

Neuroscience cluster	Incidence in neuroscience (%)	Relative philosophy visibility (%)	Difference (p.p.)
COGNITIVE & SYSTEMS NEUROSCIENCE	25.3	52.7	+27.5
MOLECULAR & CELLULAR NEUROSCIENCE	37.0	22.1	-14.8
NEUROLOGY & NEUROSURGERY	21.3	12.4	-8.9
NEUROPATHOLOGY	8.4	2.7	-5.7
NEUROPSYCHIATRY	8.1	10.0	+1.9

6 Discussion

In this section, we build on our analyses to propose some answers to the research questions R1-R4 and to the hypotheses H1-H2.

6.1 Research question 1: Visibility over time

The visibility of neuroscience in philosophy has progressively increased (triplicated) from 1980 to 2020 (Fig. 2A). Was this increase driven by an interest specific to philosophers, or either by the general increase in the incidence of neuroscience vis-à-vis other scientific fields? (Fig. 2C) Indeed, during the last two decades of the Twentieth century, neuroscience was surrounded by a halo of hype (Cooper & Shallice, 2010); a hype that coalesced into big research programs such as the Decade of the Brain in US (1990–1999) and later in the Human Brain Project in EU (2013–2023); as well as in a steep increase in the incidence of neuroscience in WoS until 2000. However, the visibility of neuroscience in philosophy increases significantly during the very decades where the incidence of neuroscience decreases (2000–2020).

This mismatch is consistent with two views. On the former view, the visibility of neuroscience in philosophy behaves like a *salmon*: it moves upward, against the descending tide. This view may be consistent with the idea that, while the study of the mind was undergoing what has been called “a cognitive neuroscience revolution” (Boone & Piccinini, 2016), philosophers took interest in the heated debates surrounding the epistemological appraisal of new methods and their shortcomings (cf. Coltheart, 2006; Roskies, 2009). On the second view, the increased visibility of neuroscience in philosophy is an aftermath of the development of neuroscientific research, with a twenty-years delay. Given that philosophical papers are usually longer than neuroscientific ones, this delay may be partly due to the longer time they take to be published and circulated. Hence, philosophy acts toward neuroscience like the Hegelian *owl of Minerva* which spreads its wings only with the falling of dusk.

This view would entail that philosophers' citations to neuroscience journals (in our terms, the visibility of neuroscience in philosophy) would be mainly post-hoc reflections of the philosophical implications of neuroscientific findings. These two views are not mutually exclusive, and may refer to different areas of philosophy (more on this below, in Sect. 6.2).

On the other hand, the visibility of philosophy in neuroscience continues to grow (Fig. 2B) irrespective of the incidence of philosophy on the total (albeit remaining an order of magnitude smaller than the converse; Fig. 2D), suggesting that some channels for scientific communication have been successfully established.

In any case, noting that the increase in visibility is stable and affects both fields – despite the different size, – it looks like, 45 years after the Sloan Report, the intersection between philosophy and neuroscience has finally become “the focus of formally recognized scholarly effort” (but see below, Sect. 6.6).

6.2 Research question 2: Neuroscience in philosophy

As we observed, the visibility of neuroscience in philosophy varies greatly both between clusters. It is worth noting that the increase in visibility is largely driven by some specific clusters, while remaining negligible in others (cf. Fig. 3; Table 2).

Even within most of these neuro-friendly clusters, the visibility of neuroscience is not uniform. In APPLIED ETHICS and HISTORY & SOCIOLOGY OF SCIENCE, neuroscience is very visible only in those journals that have a dedicated neuroscientific focus, such as the *Journal of the History of the Neurosciences* and in *Neuroethics*.¹¹ Whereas it is virtually invisible in *Historia Mathematica* or *Nursing Ethics*. On the contrary, within PHILOSOPHY OF SCIENCE, neuroscience remains visible across most journals rather than being confined to specific “ghetto-journals”. This may indicate that philosophy of neuroscience has managed to become a recognized sub-discipline of philosophy of science (or, more properly, a sub-discipline of philosophy of psychology, as we argue in Sect. 6.6).

The high visibility of neuroscience in PHILOSOPHY OF MIND journals may seem obvious. But consider that, until the Eighties, the dominant view on the mind held that the identity of mental state (or psychological predicates) is independent by their neural realizers (Fodor, 1974). At the very least, the increased visibility of neuroscience within philosophy of mind suggests that the dominant view is no longer hegemonic.

Coming back to the animal metaphors from the previous sub-section, we suspect that the visibility of neuroscience within PHILOSOPHY OF SCIENCE is due to a salmon-like behavior, i.e., the clusters are developing an interest for the epistemological issues in neuroscience, irrespective of the (relative) deflation of the incidence of neuroscientific publications over publications in general. Whereas the visibility of neuroscience within other fields, especially APPLIED ETHICS, reflects an Owl of Minerva

¹¹ Note that the *Journal of the History of the Neurosciences* is the only journal that belongs both to NEU_N and PHI_N . This high visibility of neuroscience in this journal could be due therefore to the citations given by this journal to itself (journal self-citations).

behavior, i.e., philosophers tackling issues raised by neuroscientific findings with some delay (as suggested also by the topics most often debated at the International Society for Neuroethics from 2011 to 2020; Wei et al., 2024). If our intuition is right, in the next decade or so we can expect that the visibility of neuroscience will remain more or less stable in PHILOSOPHY OF SCIENCE, but slowly decrease in APPLIED ETHICS and other clusters – whose attention may be possibly driven away by the hype surrounding Artificial Intelligence.

What about PHILOSOPHY OF MIND? Here, the fine-grained citation analysis of a sample of philosophy of mind articles citing scientific articles performed by Yan and Liao (2023) suggests that, when philosophers cite science, many citations have merely perfunctory function (i.e., contextualizing a debate or showing off encyclopedic knowledge). However, in many other cases, citations are made for the sake of “theory-tinkering”, i.e., revising or objecting a scientific theory, or extending its scope – an especially prominent case being the fervid philosophical discussion sprouted from Karl Friston’s free energy principle (Friston, 2010; cf. Petrovich & Viola, 2024). While our PHILOSOPHY OF MIND cluster does not fully overlap with Yan and Liao’s corpus, we deem possible that the high visibility and inverse visibility of PHILOSOPHY OF MIND may reflect an active engagement in debating scientific theorizing.

6.3 Research question 3: Philosophy in neuroscience

Just like not neuroscience is not equally visible in all philosophical areas (in some areas, it is not visible at all), the visibility of philosophy is uneven across neuroscientific clusters. In particular, philosophy is significantly more visible in two clusters, COGNITIVE AND SYSTEMS NEUROSCIENCE and NEUROPSYCHIATRY (Fig. 4; Table 3). Why? A simple explanation regarding COGNITIVE AND SYSTEMS NEUROSCIENCE is that they are closely intertwined with experimental psychology. Experimental psychology and philosophy are close relatives: after all, according to most textbooks, Wundt, a professor of philosophy, founded the first psychology laboratory in Leipzig in 1879. Moreover, in his troubled history, experimental psychology often had to profoundly rethink its methods (Danziger, 1994) – a trend that might be still ongoing – and reinventing methods is something that often calls more philosophical efforts that simply working within the consolidated rails of what Kuhn (1962/1969) would call “normal science”.

Similarly, the (relatively) high visibility of philosophy in NEUROPSYCHIATRY may be due to some epistemological discussions. This speculation may be corroborated by glancing at the temporal development of the visibility of philosophy in NEUROPSYCHIATRY (Supplementary Materials, Figure S2), which increases especially in the decade 2010–2020. In fact, in 2013 the publication of the fifth edition of the Diagnostic-Statistical Manual (DSM-V), the most widespread handbook for classifying psychiatric disorders, has likely been anticipated and followed by reflections on the ontology of psychiatry. On top of that, the very nature of neuropsychiatry constantly requires to weigh neurobiological and ethical considerations. It is thus unsurprising

that neuropsychiatry was among the most prominent topics of the publications classified as ‘neuroethics’ by Leefman et al. (2016) in the 1995–2012 period, and remains a prominent topic also in the abstract presented at the International Society for Neuroethics from 2011 to 2020 (Wei et al., 2024).

6.4 Research question 4: The trading zone

As we have seen, the trading zone in philosophy is characterized by a series of journals whose visibility and inverse visibility in neuroscience are both relatively high, such as *Philosophical Psychology*, *Biology & Philosophy*, and *Neuroethics* (Fig. 5A). In general, the visibility of neuroscience within a philosophical journal exhibits a linear correlation with its inverse visibility. This suggests that philosophical journals that pay attention to neuroscientific discussions are reciprocated. As discussed above, this may reflect that some philosophical journals, especially in the PHILOSOPHY OF MIND cluster, seek to participate to the debate of some specific neuroscientific theory (as shown by the visibility therein; see Supplementary Materials, Figure S1); and somehow succeeds in so doing (as reflected by their inverse visibility; see Table 2).

Instead, rather than a trading zone with a bidirectional citation flow, the citation relationship of neuroscience journals with respect to philosophy loosely recalls the organization of neurons: on the one hand, we find axon-like journals, that give citation but do not receive them (Fig. 5B). Here, the visibility of philosophy is high (as compared to other neuroscientific journals), but the inverse visibility is not. Journals included here are usually explicitly interdisciplinary and open to theoretical contributions, such as *Cognitive Systems Research*, or *NeuroQuantology* (the journal description explicitly states that it “aims to promote scientific dialogue and collaboration among researchers, practitioners, and scholars in the fields of neuroscience, quantum physics, psychology, philosophy, and related disciplines”¹²).

On the other hand, there are dendrite-like journals, which receive many citations from philosophy without reciprocating. In other terms, the visibility of philosophy here is pretty low, but the inverse visibility is pretty high. The two neuroscientific journals for which the inverse visibility in philosophy is higher (despite philosophy being poorly visible) are *Trends in Cognitive Science* and *Nature Reviews Neuroscience* (see also Table 4). Why they are often cited by philosophers should be relevantly easy to explain. First, both are popular journals in general, even within their own field. Second, rather than new experiments, they mostly publish review articles. Since this kind of articles afford a bird-eye view of some areas of neuroscientific literature without zooming into too many technicalities, they might be more appealing to a readership from outside the field. Yan and Liao’s (2023) fine-grained analysis of the scientific papers most cited in “empirically informed philosophy of mind” also shows that review or theoretical articles (as opposed to new experiments) are by large the most cited (and often misinterpreted, they claim).

¹² <https://www.neuroquantology.com/>

6.5 Hypothesis H1: A neuro-revolution has occurred in philosophy

Has the revolution called for by Patricia Churchland (1986) occurred, or has it “fell”, as later noted by Bickle (2019)? The question hardly admits a straightforward answer. On the one hand, the constant and bi-directional increase in visibility between the two disciplines from 2000 to 2020 suggests that the revolution might well be occurred (Fig. 2A, B). On the other hand, it may be argued that a constant linear increase spanning 20 years signals a successful reform rather than a revolution, for which we would expect an abrupt increase concentrated in a single decade. Moreover, we ought to recall that Churchland’s agenda did not imply the development of specific areas where philosophy engaged with neuroscience, but rather a pervasive increase of visibility of neuroscience all over the board. This is not what we find, since the visibility of neuroscience has increased only slightly (or not at all) in many areas of philosophy (Table 2). Most notably, it has poorly affected the cluster of ANALYTIC PHILOSOPHY, which includes several prestigious and jealously gatekept journals with the power of changing the career of philosophers in the Anglo-American area (Katzav & Vaesen, 2017a, 2017b).

Moreover, it is worth recalling that Churchland’s agenda not only implied that philosophy ought to be reduced to neuroscience. Instead, it is more correctly construed as philosophy and neuroscience merging into each other. And while the visibility of philosophy in neuroscience increased indeed (Fig. 2B), it remains so low (about 1,5%) that unification seems nowhere in sight, even at its peak (1,8% in COGNITIVE AND SYSTEMS NEUROSCIENCE; see Table 3). In the light of this, we are inclined to deny that a full-fledged revolution has occurred – although we note that a successful reform has arisen from its ashes.

6.6 Hypothesis H2: Psychology is the disciplinary link between philosophy and neuroscience

Not all areas of neuroscience are equally interesting for, nor interested in, philosophy. As noted by Bickle (2003b), most philosophers claiming to be interested in neuroscience are actually interested in cognitive neuroscience. As we have shown, this interest does not merely reflect the incidence of COGNITIVE AND SYSTEMS NEUROSCIENCE (Table 5). This relative neglect of biomedical neuroscience can have theoretical repercussion. Haueis (2018) has warned against the risks of cognitive myopia, i.e. failing to consider the biological functions of brain structures and processes (e.g., homeostasis), (mis)interpreting them in psychological terms instead. After all, before being an organ that supports cognition (and in order to do so), the brain is a biological organ, with its own cytology, histology, and metabolism.

Anyhow, the skewedness in visibility is symmetrical. The clusters of neuroscience (though not necessarily the journals) where philosophy is more visible are the same areas that are more visible in philosophy (Table 4). In sum, it seems reasonable to affirm that psychology is the driving belt between philosophy and neuroscience. Even if half a century after the Sloan Report the connection between the

two disciplines has managed to become “the focus of formally recognized scholarly effort”, this connection is largely mediated by a middle-ground, i.e., psychology.

While we do not want to draw strong axiological conclusions from this fact, we cannot help wondering what would result from the encounter between philosophy and molecular neuroscience or neurosurgery, a largely unexplored field thus far.

7 Conclusions

While at the end of the Seventies philosophy and neuroscience were distant fields, the growing number of conferences, publications, and other initiatives focusing on linking them together suggests that the two fields are approaching.

In this study, we provided a quantitative bird-eye portrayal of the philosophy-neuroscience nexus using advanced methods from citation analysis. Combining suitable bibliometric indicators of cross-field visibility with journal citation mapping techniques, we were able to investigate four dimensions of the nexus: how the visibility of neuroscience in philosophy and of philosophy in neuroscience has changed over time, in which areas of philosophy neuroscience is more visible, in which areas of neuroscience philosophy is more visible, and how the trading zone between the two fields is configured. We also discussed the extent to which evidence supports two specific hypotheses about the relationship between philosophy and neuroscience: the supposed occurrence of a neuro-revolution in philosophy and the role of psychology as the disciplinary link between neuroscience and philosophy.

Our results show that both the visibility of neuroscience in philosophy and the visibility of philosophy in neuroscience have increased significantly since the turn of the century, albeit the latter remains an order of magnitude lower than the former. Neuroscience, however, has not had the same impact on all areas of philosophy. The subdisciplines where neuroscientific literature is more visible are philosophy of mind, applied ethics, and philosophy of science, whereas core areas in (analytic) philosophy have remained relatively isolated from neuroscience. Similarly, the visibility of philosophy in neuroscience is not distributed homogeneously over the field but is concentrated in cognitive and systems neuroscience and neuropsychiatry. The impact of philosophy in biomedical areas, by contrast, is virtually nonexistent. As to the trading zone between philosophy and neuroscience, we observed that it works differently in philosophy and in neuroscience. In philosophy, the same journals where neuroscience is highly visible are themselves quite visible in neuroscience, showing that they function as active loci of bidirectional citation flows. In neuroscience, on the other hand, journals where philosophy is particularly visible are not themselves particularly visible in philosophy, and vice versa, suggesting that the loci of import of philosophy into neuroscience do not coincide with the loci of export of neuroscience to philosophy.

Regarding our two hypotheses, data do not support the hypothesis that a widespread neuro-revolution has transformed philosophy radically, as the impact of neuroscience has been localized in specific areas rather than affecting the whole field. Data, however, support the hypothesis that psychology functions as a mediating disciplinary link between philosophy and neuroscience, as cognitive

neuroscience has a disproportionately higher visibility in philosophy compared to biomedical areas and, at the same time, cognitive neuroscience is the area where philosophy is more visible.

These results show that our methodology based on bibliometric indicators and science mapping is suitable for revealing several features of field-size phenomena of interaction between disciplines, which escape the traditional methods of philosophy of science. We note, however, that other methodological choices could be made in order to investigate the same phenomena from further angles. In particular, we based our bibliometric delineation of philosophy and neuroscience on the journal-level classification system of Web of Science. However, in some cases, neuroscientists publish in philosophy journals and philosophers in neuroscience ones, blurring the one-to-one correspondence between philosophy/neuroscience journals and philosophy/neuroscience authors (as we detail in Petrovich & Viola, 2024). Such cross-field publication activity at the author level is not captured by our method but offers an additional perspective to investigate quantitatively the knowledge exchange between the two fields. In fact, depending on how disciplines are bibliometrically operationalized, several different measures of interdisciplinary communication can be devised (Porter et al., 2006). In this study, we favored an *information flow* approach based on citation exchanges (Leydesdorff & Rafols, 2011), but alternative approaches can be pursued. For instance, an organizational approach based on researchers' affiliations and collaboration patterns is possible (see e.g., Abramo et al., 2012), as is a content-based approach based on similarities of words or topics contained in the full text of articles (see e.g., Malaterre et al., 2019). Future research could focus also on further actors of the publication system, such as journal editors (c.f. Baccini & Barabesi, 2010). Intra-disciplinary, a heavily gate-kept journal system may contribute to attenuate the visibility of other fields. Interdisciplinary, investigating the cross-presence of philosophers in neuroscience journal boards and neuroscientists in philosophy journal boards could indeed illuminate the processes of gate-keeping that govern disciplinary boundaries.

This study confirms that quantitative analysis is a valuable tool to improve our understanding of how philosophy interacts with other disciplines. Among other things, we showed that the richness of the dynamics of the philosophy-neuroscience nexus cannot be satisfactorily captured by simplistic labels such as “revolution” but require careful empirical analyses. We hope that by gaining a better understanding how philosophy communicates with neuroscience, we can also understand each of them better.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s13194-024-00621-5>.

Acknowledgements We are grateful to Marco Facchin and Fabrizio Calzavarini for their insightful comments. Eugenio Petrovich is grateful to the Centre for Science and Technology Studies (CWTS) at Leiden University for hosting him as a guest researcher and giving him access to the CWTS database system. He gratefully acknowledges financial support from Italian Ministry of University and Research, under PNRR Young Researchers initiative (grant agreement n° SOE_0000149).

Author contribution Eugenio Petrovich: Conceptualization, Data curation, Formal analysis, Methodology, Visualization, Writing—original, Writing—review and editing. Marco Viola: Conceptualization, Methodology, Data curation, Writing—original, Writing—review and editing.

Funding Eugenio Petrovich acknowledges financial support from Italian Ministry of University and Research, under PNRR Young Researchers initiative (grant agreement n° SOE_0000149).

Data availability The raw data used in the study cannot be made available as they are proprietary data of Clarivate Analytics. All the aggregated data used in the study in the form of maps and bibliometric statistics, however, are available in the Supplementary Materials.

Declarations

Conflict of interest N/A.

Ethical approval N/A.

Informed consent N/A.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abramo, G., D'Angelo, C. A., & Di Costa, F. (2012). Identifying interdisciplinarity through the disciplinary classification of coauthors of scientific publications. *Journal of the American Society for Information Science and Technology*, 63(11), 2206–2222. <https://doi.org/10.1002/asi.22647>
- Allen, C., & Beavers, T. (2011). Synthese special issue: Representing philosophy. *Synthese*, 182(2), 181–183. <https://doi.org/10.1007/s11229-009-9664-z>
- Anderson, M. L. (2014). *After phrenology* (Vol. 547). MIT Press.
- Baccini, A., & Barabesi, L. (2010). Interlocking editorship. A network analysis of the links between economic journals. *Scientometrics*, 82(2), 365–389. <https://doi.org/10.1007/s11192-009-0053-7>
- Baldwin, M. C. (2015). *Making Nature: The history of a scientific journal*. University of Chicago Press.
- Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. Psychology Press.
- Bickle, J. (2003a). Editor's note. *Brain and Mind*, 4, 305.
- Bickle, J. (2003b). *Philosophy and neuroscience: A ruthlessly reductive account*. Springer Science & Business Media.
- Bickle, J. (2019). Lessons for experimental philosophy from the rise and “fall” of neurophilosophy. *Philosophical Psychology*, 32(1), 1–22.
- Bickle, J., Mandik, P., & Landreth, A. (1999/2019). The philosophy of neuroscience. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/fall2019/entries/neuroscience/>. Accessed 01 April 2024.
- Bickle, J., Einstein, G., & Hardcastle, V. (2000). Editors' Introduction. *Brain and Mind*. Springer.
- Bonino, G., Maffezoli, P., Petrovich, E., & Tripodi, P. (2022). When philosophy (of science) meets formal methods: A citation analysis of early approaches between research fields. *Synthese*, 200(2), 177. <https://doi.org/10.1007/s11229-022-03484-6>
- Boone, W., & Piccinini, G. (2016). The cognitive neuroscience revolution. *Synthese*, 193, 1509–1534.
- Börner, K. (2010). *Atlas of science: Visualizing what we know*. MIT Press.
- Boyack, K. W., & Klavans, R. (2019). Creation and analysis of large-scale bibliometric networks. In W. Glänzel, H. F. Moed, U. Schmoch, & M. Thelwall (Eds.), *Springer Handbook of Science and*

- Technology Indicators* (pp. 187–212). Springer International Publishing. https://doi.org/10.1007/978-3-030-02511-3_8
- Calzavarini, F., & Cevolani, G. (2022). Abductive reasoning in cognitive neuroscience: Weak and strong reverse inference. *Synthese*, 200(2), 70.
- Chi, P.-S., & Conix, S. (2022). Measuring the isolation of research topics in philosophy. *Scientometrics*, 127(4), 1669–1696. <https://doi.org/10.1007/s11192-022-04276-y>
- Churchland, P. S. (1986). *Neurophilosophy*. MIT Press.
- Churchland, P. S. (1987). Epistemology in the age of neuroscience. *Journal of Philosophy*, 84(10), 544–553.
- Coltheart, M. (2006). What has functional neuroimaging told us about the mind (so far)? *Cortex; a Journal Devoted to the Study of the Nervous System and Behavior*, 42(3), 323–331.
- Cooper, R. P., & Shallice, T. (2010). Cognitive neuroscience: The troubled marriage of cognitive science and neuroscience. *Topics in Cognitive Science*, 2(3), 398–406.
- Craver, C. F. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. Oxford : Clarendon Press.
- Csiszar, A. (2018). *The scientific journal: Authorship and the politics of knowledge in the nineteenth century*. The University of Chicago Press.
- Danziger, K. (1994). *Constructing the subject: Historical origins of psychological research*. Cambridge University Press.
- Fodor, J. A. (1974). Special sciences (or: The disunity of science as a working hypothesis). *Synthese*, 28(2), 97–115.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
- Galison, P. (1997). *Image and logic: A material culture of microphysics*. University of Chicago Press.
- Gardner, H. (1987). *The mind's new science: A history of the cognitive revolution*. Basic Books.
- Hauéis, P. (2018). Beyond cognitive myopia: A patchwork approach to the concept of neural function. *Synthese*, 195(12), 5373–5402.
- Higgins, A., & Smith, B. (2013). A citation based view of the ontology community in philosophy. In H. Davis (Ed.), *Proceedings of the 5th annual ACM web science conference*. ACM Digital Library, Association for Computing Machinery.
- Hyland, K. (1999). Academic attribution: Citation and the construction of disciplinary knowledge. *Applied Linguistics*, 20(3), 341–367. <https://doi.org/10.1093/applin/20.3.341>
- Janssen, A., Klein, C., & Slors, M. (2017). What is a cognitive ontology, anyway? *Philosophical Explorations*, 20(2), 123–128.
- Katzav, J., & Vaesen, K. (2017a). On the emergence of American analytic philosophy. *British Journal for the History of Philosophy*, 25(4), 772–798. <https://doi.org/10.1080/09608788.2016.1261794>
- Katzav, J., & Vaesen, K. (2017b). Pluralism and Peer Review in Philosophy. *Philosophers' Imprint*, 17(19), 1–20.
- Kessler, M. M. (1963). Bibliographic coupling extended in time: Ten case histories. *Information Storage and Retrieval*, 1(4), 169–187. [https://doi.org/10.1016/0020-0271\(63\)90016-0](https://doi.org/10.1016/0020-0271(63)90016-0)
- Khelifaoui, M., Gingras, Y., Lemoine, M., & Pradeu, T. (2021). The visibility of philosophy of science in the sciences, 1980–2018. *Synthese*. <https://doi.org/10.1007/s11229-021-03067-x>
- Kitcher, P. (2011). Philosophy inside out. *Metaphilosophy*, 42(3), 248–260. <https://doi.org/10.1111/j.1467-9973.2011.01684.x>
- Knobe, J. (2015). Philosophers are doing something different now: Quantitative data. *Cognition*, 135, 36–38.
- Kostić, D., & Halfman, W. (2023). Mapping explanatory language in neuroscience. *Synthese*, 202(4), 112.
- Kreuzman, H. (2001). A co-citation analysis of representative authors in philosophy: Examining the relationship between epistemologists and philosophers of science. *Scientometrics*, 51(3), 525–539.
- Kuhn, T. S. (1962/1969). *The structure of scientific revolutions*. University of Chicago press.
- Larivière, V., Archambault, É., Gingras, Y., & Vignola-Gagné, É. (2006). The place of serials in referencing practices: Comparing natural sciences and engineering with social sciences and humanities. *Journal of the American Society for Information Science and Technology*, 57(8), 997–1004. <https://doi.org/10.1002/asi.20349>
- Leeffmann, J., Levallois, C., & Hilt, E. (2016). Neuroethics 1995–2012. A bibliometric analysis of the guiding themes of an emerging research field. *Frontiers in Human Neuroscience*, 10, 336.

- Leydesdorff, L., & Rafols, I. (2011). Indicators of the interdisciplinarity of journals: Diversity, centrality, and citations. *Journal of Informetrics*, 5(1), 87–100. <https://doi.org/10.1016/j.joi.2010.09.002>
- Leydesdorff, L., Hammarfelt, B., & Salah, A. (2011). The structure of the Arts & Humanities Citation Index: A mapping on the basis of aggregated citations among 1,157 journals. *Journal of the American Society for Information Science and Technology*, 62(12), 2414–2426. <https://doi.org/10.1002/asi.21636>
- Malaterre, C., Chartier, J.-F., & Pulizzotto, D. (2019). What is this thing called philosophy of science? A computational topic-modeling perspective, 1934–2015. *HOPOS: The Journal of the International Society for the History of Philosophy of Science*, 9(2), 215–249.
- McCaffrey, J. B. (2023). Evolving Concepts of Functional Localization. *Philosophy Compass*, 18(5), e12914.
- McLevey, J., Graham, A. V., McIlroy-Young, R., Browne, P., & Plaisance, K. S. (2018). Interdisciplinarity and insularity in the diffusion of knowledge: An analysis of disciplinary boundaries between philosophy of science and the sciences. *Scientometrics*, 117(1), 331–349. <https://doi.org/10.1007/s11192-018-2866-8>
- Mingers, J., & Leydesdorff, L. (2015). A review of theory and practice in scientometrics. *European Journal of Operational Research*, 246(1), 1–19. <https://doi.org/10.1016/j.ejor.2015.04.002>
- Moravcsik, M. J., & Murugesan, P. (1975). Some results on the function and quality of citations. *Social Studies of Science*, 5(1), 86–92. <https://doi.org/10.1177/030631277500500106>
- Nathan, M. J., & Del Pinal, G. (2017). The future of cognitive neuroscience? *Reverse Inference in Focus. Philosophy Compass*, 12(7), e12427.
- Noichl, M. (2019). Modeling the structure of recent philosophy. *Synthese*. <https://doi.org/10.1007/s11229-019-02390-8>
- Pence, C. H., & Ramsey, G. (2018). How to do digital philosophy of science. *Philosophy of Science*, 85(5), 930–941. <https://doi.org/10.1086/699697>
- Petrovich, E. (2018). Accumulation of knowledge in para-scientific areas: The case of analytic philosophy. *Scientometrics*, 116(2), 1123–1151. <https://doi.org/10.1007/s11192-018-2796-5>
- Petrovich, E. (2021). Science Mapping and Science Maps. *Knowledge Organization*, 48(7–8), 535–562. <https://doi.org/10.5771/0943-7444-2021-7-8-535>
- Petrovich, E. (2024). A quantitative portrait of analytic philosophy: Looking through the margins. *Springer International Publishing*. <https://doi.org/10.1007/978-3-031-53200-9>
- Petrovich, E., & Buonomo, V. (2018). Reconstructing late analytic philosophy. *A Quantitative Approach. Philosophical Inquiries*, 6(1), 151–181. <https://doi.org/10.4454/philiinq.v6i1.184>
- Petrovich, E., & Viola, M. (2024). Dalla filosofia alle neuroscienze (passando per la psicologia?). *Una Prospettiva Scientometrica. Sistemi Intelligenti*, 1, 47–72. <https://doi.org/10.1422/112959>
- Porter, A. L., Roessner, J. D., Cohen, A. S., & Perreault, M. (2006). Interdisciplinary research: Meaning, metrics and nurture. *Research Evaluation*, 15(3), 187–196. <https://doi.org/10.3152/147154406781775841>
- Poldrack R. A. (2006). Can cognitive processes be inferred from neuroimaging data?. *Trends in Cognitive Sciences*, 10(2), 59–63. <https://doi.org/10.1016/j.tics.2005.12.004>
- Pradeu, T., Lemoine, M., Khelifaoui, M., & Gingras, Y. (2021). Philosophy in Science: Can philosophers of science permeate through science and produce scientific knowledge? *The British Journal for the Philosophy of Science*, 715518. <https://doi.org/10.1086/715518>
- Price, C. J., & Friston, K. J. (2005). Functional ontologies for cognition: The systematic definition of structure and function. *Cognitive Neuropsychology*, 22(3–4), 262–275.
- Ramsey, G., & Block, A. de (Eds.). (2022). *The dynamics of science: Computational frontiers in history and philosophy of science*. University of Pittsburgh Press.
- Rescher, N. (2005). The taxonomy of philosophy and its fate. In *Collected papers. 1: Studies in 20th century philosophy* (pp. 61–87). Ontos Verlag.
- Roskies, A. (2002). Neuroethics for the new millenium. *Neuron*, 35(1), 21–23.
- Roskies, A. L. (2009). Brain-mind and structure-function relationships: A methodological response to Coltheart. *Philosophy of Science*, 76(5), 927–939.
- Salton, G., & McGill, M. J. (1983). *Introduction to modern information retrieval*. McGraw-Hill.
- Small, H. (2003). Paradigms, citations, and maps of science: A personal history. *Journal of the American Society for Information Science and Technology*, 54(5), 394–399. <https://doi.org/10.1002/asi.10225>

- State of the Art Committee (1978). *Cognitive Science, 1978. Report of the state of the art committee to the advisors of the Alfred P. Sloan foundation*. Available at the URL https://digitalcollections.library.cmu.edu/node/31966?search_api_fulltext. Accessed 01 April 2024.
- Sugimoto, C. R., & Larivière, V. (2018). *Measuring research: What everyone needs to know*. Oxford University Press.
- Sugimoto, C. R., & Weingart, S. (2015). The kaleidoscope of disciplinarity. *Journal of Documentation, 71*(4), 775–794. <https://doi.org/10.1108/JD-06-2014-0082>
- van Eck, N. J., & Waltman, L. (2009). How to normalize cooccurrence data? An analysis of some well-known similarity measures. *Journal of the American Society for Information Science and Technology, 60*(8), 1635–1651. <https://doi.org/10.1002/asi.21075>
- van Eck, N. J., & Waltman, L. (2010). Software survey: VOSviewer, a computer program for bibliometric mapping. *Scientometrics, 84*(2), 523–538. <https://doi.org/10.1007/s11192-009-0146-3>
- van Eck, N. J., Waltman, L., Dekker, R., & van den Berg, J. (2010). A comparison of two techniques for bibliometric mapping: Multidimensional scaling and VOS. *Journal of the American Society for Information Science and Technology, 61*(12), 2405–2416. <https://doi.org/10.1002/asi.21421>
- Van Raan, A. F. J. (1998). In matters of quantitative studies of science the fault of theorists is offering too little and asking too much. *Scientometrics, 43*(1), 129–139. <https://doi.org/10.1007/BF02458401>
- Waltman, L., van Eck, N. J., & Noyons, E. C. M. (2010). A unified approach to mapping and clustering of bibliometric networks. *Journal of Informetrics, 4*(4), 629–635. <https://doi.org/10.1016/j.joi.2010.07.002>
- Waltman, L., & van Eck, N. J. (2014). Visualizing bibliometric networks. In Y. Ding, R. Rousseau, & D. Wolfram (Eds.), *Measuring scholarly impact: Methods and practice* (pp. 285–320). Springer.
- Wei, N. Y., Choi, R. J., Sullivan, L. S., & Wexler, A. (2024). Who does neuroethics scholarship address, and what does it recommend? A content analysis of selected abstracts from the international neuroethics society annual meetings. *Neuroethics, 17*(2), 1–10.
- Yan, K., & Liao, C. Y. (2023). A co-citation analysis of cross-disciplinarity in the empirically-informed philosophy of mind. *Synthese, 201*(5), 159.
- Zeki, S. (1999). *Inner vision: An exploration of art and the brain*. Oxford University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.