



Epistemic integration and social segregation of AI in neuroscience

Sylvain Fontaine, Floriana Gargiulo, Michel Dubois, Paola Tubaro

► To cite this version:

Sylvain Fontaine, Floriana Gargiulo, Michel Dubois, Paola Tubaro. Epistemic integration and social segregation of AI in neuroscience. 2023. <hal-04225888>

HAL Id: hal-04225888

<https://hal.science/hal-04225888v1>

Preprint submitted on 3 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Epistemic integration and social segregation of AI in neuroscience

Sylvain Fontaine^{1*}, Floriana Gargiulo¹, Michel Dubois¹,
Paola Tubaro²

¹GEMASS, CNRS-Sorbonne Université, 59-61 rue Pouchet, 75017 Paris,
France.

²CREST, CNRS-IPP, 5 avenue Henry Le Chatelier, 91120 Palaiseau,
France.

*Corresponding author(s). E-mail(s): syl20.fontaine@outlook.fr;

Abstract

In recent years, Artificial Intelligence (AI) shows a spectacular ability of insertion inside a variety of disciplines which use it for scientific advancements and which sometimes improve it for their conceptual and methodological needs. According to the transverse science framework originally conceived by Shinn and Joerges, AI can be seen as an instrument which is progressively acquiring an universal character through its diffusion across science. In this paper we address empirically one aspect of this diffusion, namely the penetration of AI into a specific field of research. Taking neuroscience as a case study, we conduct a scientometric analysis of the development of AI in this field. We especially study the temporal egocentric citation network around the articles included in this literature, their represented journals and their authors linked together by a temporal collaboration network. We find that AI is driving the constitution of a particular disciplinary ecosystem in neuroscience which is distinct from other subfields when regarding the references, and which is gathering atypical scientific profiles who are coming from neuroscience or outside it. Moreover we observe that this AI community in neuroscience is socially confined in a specific zone of the neuroscience collaboration network, which is also keeping to publish in a small set of dedicated journals that are mostly active in AI research. According to these results, the diffusion of AI in a discipline such as neuroscience didn't really challenge its disciplinary orientations but rather induced the constitution of a dedicated socio-cognitive workforce inside this field.

Keywords: Science of science, Artificial Intelligence, neuroscience, research-technology

1 Introduction

In recent years, Artificial Intelligence (AI) is continuously spreading across science. The associated worldwide scientific production and the amount of dedicated funding programs for technological developments supported both by academia and industry, are spectacularly growing (Baruffaldi et al., 2020; Gao & Wang, 2023; Liu, Shapira, & Yue, 2021), and the output of such research are touching a variety of disciplines that are more and more citing it (Arencibia-Jorge, Vega-Almeida, Jiménez Andrade, & Carrillo-Calvet, 2022; Frank, Wang, Cebrian, & Rahwan, 2019). Developed and confined within STEM disciplines between the late 1980's and 2010, mainly represented by mathematics, statistics and computer science, AI research is now opening to fields that are cognitively and methodologically distant from these originating ones (Gargiulo, Fontaine, Dubois, & Tubaro, 2022).

AI knowledge and tools are thus amenable to applications in such welcoming disciplines which use them for scientific advancements, and which sometimes adapt and improve them in turn for their conceptual and methodological needs. Bianchini et al. (2022) studied especially such mutual improvement dynamics at a macroscopic scale of science by focusing on the diffusion of neural networks' concepts and tools in numerous disciplines. Qualified as a *general method of invention* by these authors because of its non-disciplinary status, AI is thus expanding the *adjacent possible* of the purposes of the concerned disciplines, and is therefore the source of innovation in the latter (Kauffman, 2000; Monechi, Ruiz-Serrano, Tria, & Loreto, 2017).

In this paper we address the process of penetration of AI inside a given field of research, which is not part of originating STEM disciplines while largely applying AI. We focus on neuroscience, which are part of such fields according to Gargiulo et al. (2022). We choose especially this field as it has a particular history which is closely intertwined with the AI one. Indeed neuroscience and AI are continuously providing feed-backs to each other towards the comprehension of the human brain and the artificial reproduction of some cognitive processes (Hassabis, Kumaran, Summerfield, & Botvinick, 2017; Savage, 2019). This lead us to also assume that neuroscience, under the assumption of an AI-receiving field of research, contribute also to reshape AI knowledge, tools and practices, thus expanding the *adjacent possible* of AI.

Along these lines, we focus on two dynamical mechanisms underlying the appropriation of AI by neuroscience and then the possible knowledge transfer between the two fields:

- The cognitive embedding of AI into the multidisciplinary context of neuroscience:
 1. what is the disciplinary landscape around neuroscience research which is using AI?
 2. how such research fits in the disciplinary objectives of neuroscience?
- The diffusion throughout the neuroscience community:
 3. who are the actors leading AI research in this field?
 4. do AI-related knowledge widespread throughout the whole scientific community?

In the following we answer these questions by analyzing a large corpus of neuroscience publications published between 1970 and 2019, extracted with a journal-oriented query from the Microsoft Academic Knowledge Graph (MAG) (Färber, 2019). We distinguish AI-related publications from others with a dedicated keyword filter applied to their titles and abstracts, as in (Gargiulo et al., 2022). After discussing in Section 2 the intertwined history of AI and neuroscience, and also the main concepts mobilized in this paper, we detail in Section 3 the building of some relational structures and some metrics given by the corpus that will be analyzed afterwards, namely its temporal egocentric citation network (composed of bibliographical references and citations of each paper), its peer-reviewed journals, and its temporal collaboration network, featured with the disciplinary profiles of the scientists.

In Sections 4.1 and 4.2, we identify within the citation network the disciplines that are shaping the AI research in neuroscience, and to what extent they differs from those that are characterizing other neuroscience’s subfields that are not concerned with AI. While the influences of the AI research in neuroscience are progressively shifting since 1970 from biomedical science to the STEM disciplines, and is building its own knowledge basis inside neuroscience, it keeps to fit in the disciplinary objectives of the latter. We also confirm this partial cognitive integration in Section 4.3, which reveals the spreading of AI research in almost all the journals of the database, but with different patterns of promotion of such research. In particular, we show that the journals that are publishing a lot of AI-related works are invested by researchers who contributed to AI research in general (inside or outside neuroscience). By regarding the temporal co-authorship network of the field, we show in the following Section 4.4 that the AI community is polarized between two main academic profiles, a biomedical and clinical one integrated inside neuroscience while being a few involved in AI, and another much more involved in AI research but remaining close to STEM disciplines. Putting all these scientists together, we then show that they are more and more isolated in the collaboration network since 1970.

Finally, by considering AI as a *research-technology* spreading across science, as originally conceived by Shinn and Joerges (2002) and formalized later by Hentschel (2015), we finally discuss in Section 5 the partial diffusion observed empirically in our results. With this framework we also discuss the similarity and differences between the local development dynamics of AI within this particular field of research and the global one within the whole science system that is depicted in (Gargiulo et al., 2022).

2 Literature review

2.1 AI and neuroscience: toward the mutual expansion of their *adjacent possible*

Artificial intelligence (AI) commonly refers to both a research program and, more generally, a set of complex computer-based programs which aim to mimic human mind processes with high reckoning power. Although its foundations are mainly associated with STEM disciplines, mainly mathematics, statistics and computer science (Gargiulo et al., 2022), the AI is constantly evolving alongside the neuroscience by maintaining a virtual circle of mutual improvement (Hassabis et al., 2017).

Indeed neuroscience bring in a first place empirical confirmation of some theoretical models that reproduce parts of mental processes, and that were first imagined, analytically derived and computationally simulated by (neuro)psychologists into the field of cognitive science (Cooper & Shallice, 2010; Lake, Ullman, Tenenbaum, & Gershman, 2017). Most of these models are at the roots of AI-related algorithms, notably the case of the bio-inspired AIs such as artificial neural networks and their numerous versions, which became *biologically plausible* with neuroscience. In particular, the 1980's and early 1990's marked also the launching of the first body-scanner machines applying positron emission tomography (PET) and functional magnetic resonance imaging (fMRI), that both enabled important discoveries on functional biological mechanisms in human brain that are induced by complex cognitive tasks (Cooper & Shallice, 2010). This contributed to the rise the connectionist paradigm that is today dominant in brain sciences, even if some debates are persisting into cognitive science about the representations of knowledge and the logical operations to process them in interaction with the real world (Andler, 1990; McCarthy, 1981; Perconti & Plebe, 2020)

In turn, in the recent *big data* era, a panel of AI tools enables today the efficient process of various biomedical data that was acquired from important clinical trials and cohorts for studying brain damages (Gopinath, 2023). They are especially part of the improvement of the diagnosis of various neuro-degenerative diseases and of the attribution of potential dedicated treatments, if they exist.

This virtual feedback loop, well documented in the neuroscience literature, is thus inducing a reinforcement dynamics of both AI and neuroscience which are receiving it. From this assertion we assume that AI and neuroscience are expanding the *adjacent possible* of each other (Kauffman, 2000; Monechi et al., 2017), ie. one is extending the field of possibilities that have yet to be explored of the other, by blending with the pre-existing knowledge and practices of the latter. One is thus reshaping the knowledge space of the other.

2.2 Conceiving AI as a *research-technology* in science

In a socio-historical perspective of science, Shinn and Joerges (2002) propose the notion of *research-technology*, called also *transverse science*, to describe the dynamics of science since the end of the 19th century, now largely based on instrumentation for experimental or empirical investigations, especially in physical and life sciences. Within such a research regime, the production of knowledge is conditioned by an instrument that is designed in a specific research environment before being disseminated outside the latter (Marcovich & Shinn, 2020; Shinn & Joerges, 2002). Such an instrument requires the involvement of a dedicated socio-cognitive workforce composed of a variety of actors (scientists, technicians, promoters, administrators,...), that is developing a dedicated technological culture of it, ie. a set of associated knowledge and practices shared by everyone in the community, without necessarily claiming together a common professional identity related to it. Regarding these criteria, they are shaping a social group that could be also defined as an *epistemic community* (Haas, 1992; Roth, 2008), although the research-technology framework relaxes the notion of socio-cognitive boundaries that are specific to such communities. According to Shinn and

Joerges (2002), the actors mentioned above are indeed able to move between different research environments and are often coming across the established disciplinary boundaries in order to provide their expertise for the resolution of diverse scientific problems.

The research-technology associated with the instrument is thus a dynamical entity with changing social and institutional organizations at different moment of its development. Based on Shinn and Joerges (2002), Hentschel (2015) proposes four stages to describe the life of a given research-technology in an historical perspective, ie. throughout long time period: *prehistory*, *exploration*, *optimization* and *diffusion*.

The first two stages are often associated together in the designing and testing phases of a given instrument. According to these authors, the underlying process needs temporarily a closure of the community working on it, thus fostering the creation of a dedicated scientific field, often interdisciplinary, in which the instrument is also a research object.

Shinn and Joerges (2002) also highlight the criterion of *genericity* of an instrument: at the later stages of its development, it could be adapted for disciplinary research contexts other than those in which it has been originally designed for, or for a variety of applications outside scientific research. The two last development steps of Hentschel are thus representing such dynamics toward this final step. In particular, the diffusion phase requires a relaxation of the disciplinary boundaries within which it has been designed during the prehistoric and exploratory phase, hence an openness of the actors in the community, which promote it in different fields of research within academia or industry. This the case of technologies related to laser beams or X-ray emission, now use widely from experimental physics to medical research and standard practices, even in daily objects such as smartphones (Hentschel, 2015).

According to the analysis of the AI research ecosystem provided by Gargiulo et al. (2022), AI seems to embrace some characteristics of an instrument embedded in a global research-technology program within science. Indeed its modern forms was designed in a bounded disciplinary environment from the 1980's until the 2010's, and it is now spreading in other disciplines. Through the analysis of the development of AI in neuroscience, although within a bigger period (1970-2019), we discuss in Section 5 to what extent this framework applies in the case of the development and the diffusion of AI in a single field of research.

3 Data and methods

3.1 Extraction and preprocessing of the data

MAG database is a big network of metadata, such as title, abstract, publication date or publisher, that are associated to a huge set of various scientific contributions published as journal articles, books or papers in conference proceedings (Färber, 2019). These metadata are embedded as strings or numerical identifiers in a set of dump files linked together by link files, all of them being available on the data deposit platform Zenodo (Färber, 2020).

To build an exhaustive database of the neuroscience literature from MAG, we draw upon a list of journals that we extracted first from the databases *Web of Science*

(WOS) and *SCImago Journal Rank* (SJR), after removing duplicates occurring in both of them. We use the SJR in addition of WOS because it provides another measure of impact that is improving the visibility of some journals that could have a lower impact factor in WOS. Moreover, because of the different disciplinary classifications provided by the databases, SJR is covering more neuroscience journals than WOS. For instance, in 2021, WOS referenced only 281 journals labeled as “neuroscience” while SJR referenced 608 journals labeled as such. We consider here the neuroscience journals referenced by WOS and other neuroscience journals referenced as such by SJR but that are labeled differently by WOS – for instance, the well-recognized journal *Neurocomputing*, while labeled with both “neuroscience” and “computer science” in SJR, is labeled only as “computer science” by WOS. We thus get 421 journals. For sake of consistency, we lie only on the disciplinary classification of WOS.

Here we focus only on peer-reviewed journals because of their easy availability in WOS database, although we are conscious that AI research could be largely published in other publication media such as preprint archives and conference proceedings (Wainer, Eckmann, Goldenstein, & Rocha, 2013).

Then we extract the publications from MAG with an ISSN identifier matching procedure based on this list of journals. These publication are brought together in a set denoted by \mathcal{P} in the following. In order to distinguish AI-related works in our neuroscience dataset, we apply a selection criterion such that they must include at least one AI-related keyword in their title or abstract (see the SI of Gargiulo et al. (2022) for the complete list of keywords). This subset is denoted $\mathcal{P} \cap AI$. All neuroscientific publications outside this set is called $\mathcal{P} \cap \overline{AI}$.

In the following we keep only the papers that have been published in the period running from 1970 until 2019, and that have at least 10 references and at least 10 generated citations. Finally the sets $\mathcal{P} \cap \overline{AI}$ and $\mathcal{P} \cap AI$ includes respectively 829,317 and 26,374 papers, spread over the period mentioned above according to Fig.1. Therefore, among all these high-impact research in neuroscience, only 3% of them contain AI-related keywords. The inset of this figure exhibits especially a slow growth of the share of AI-related publications in neuroscience at the very end of the 1980’s, which is stabilizing around 1995 and followed by a very rapid growth from 2007 until 2019. This plateau around 1995 suggests a second, prolonged period *AI winter* in neuroscience, characterized here by stable interest in AI research but not waning, unlike in other disciplines or research fields (Cardon, Cointet, & Mazières, 2018; Schuchmann, 2019). The following period of important growth of this share, started in 2007, suggests also the main influence of the rise of deep learning techniques, a well-known trend shared by almost all the sciences (Cardon et al., 2018; Gargiulo et al., 2022). The share of AI publications in neuroscience reaches only 10% of the number of publications at its highest stage situated at the end of the studied period, which means that the AI-related research remains rather marginal in neuroscience even today.




Fig. 1 Left: Classification of the papers in the extracted corpus. The cross-section between the blue and red zones is corresponding to the set $\mathcal{P} \cap \overline{AI}$ in the main text, while the blue zone excluded this cross-section is representing the set $\mathcal{P} \cap AI$. The green set $\overline{\mathcal{P}}$ includes all papers added after the building of the ego-centered citation network that are not published in neuroscience journal. Right: Cumulative number N_c of papers in the two main subsets considered in this paper. The inset shows the instantaneous part of AI-related publications in the whole neuroscience corpus.

3.2 Capturing the disciplinary landscape

3.2.1 Building of a citation ecosystem

After gathering the dataset under study, we build an exhaustive citation ecosystem around neuroscience publications, as illustrated in Fig. 2. We use the general reference link file of MAG – which is actually a giant citation network – for retrieving the direct references from and toward our identified papers. We thus add papers that are published either in neuroscience publishers, therefore in the set \mathcal{P} , or in others such that they join the set denoted by $\overline{\mathcal{P}}$, ie. the green border zone in Fig. 1. With the journal classification of the WOS, we then enrich these sets of references and citations by assigning them a set of Journal Subject Categories (JSC) associated to their journals if they are referenced in this database, which is a synonym of a large well-recognized discipline in science. The papers with no associated JSC thus are automatically discarded.

Then, for each corpus $\mathcal{P} \cap AI$ and $\mathcal{P} \cap \overline{AI}$, we count the annual number of citations obtained by each JSCs represented in the global set of bibliographic references used by their papers. Doing this year by year over the period 1970-2019, we therefore build a temporal ranking of fields of study that are most cited by each corpus. For example, in Fig. 2, the papers cited by one corpus at year y_0 (grey points) are mainly published in journals labeled by the JSC d_1 with 8 citations, followed by the JSCs d_2 with 5 citations and d_3 with only one (red points and corresponding red-contoured ranking). Hence the disciplinary ranking $r(y_0) = \{d_1, d_2, d_3\}$ with each JSC sorted according to their respective ranks mentioned above. However this ranking could be different before this year y_0 and also in the following ones, depending on the citations received by each JSC at these years.

We do the same for the JSCs that are most citing our papers of interest up to one year after their official publication year. We choose this temporal period to avoid the high time dispersion of the generated citations in the studied temporal period, the

citation impact by the papers published in 1970 being eventually higher than those published in 2019.

In this way we assess from which JSCs are issued at first the generated citations before eventually spreading into a broader disciplinary landscape. Finally, by looking at both reference and impact side, we are representing two disciplinary landscapes associated respectively with the corpora $\mathcal{P} \cap \overline{AI}$ and $\mathcal{P} \cap AI$.




Fig. 2 Citation ecosystem centered here around three papers (grey) published at year y_0 and that are included either in $\mathcal{P} \cap AI$ or $\mathcal{P} \cap \overline{AI}$. One dashed arrow represents the citation of a target paper by a source one. Hence red points are the papers that are *cited* by the paper of our corpus (*reference*) while the blue ones are *citing* them (*impact*). The rankings are shown in decreasing order with the associated number of citations of each JSC d_i .

3.2.2 Representation of the disciplinary landscape

Following the procedure described in Section 3.2.1, and according to all the couples of references' and citations' numbers associated with all the represented JSCs in $\mathcal{P} \cap \overline{AI}$ and $\mathcal{P} \cap AI$, we associate with each JSC d two couples of ranks that is indicating its weight either in references or in citations of the two corpora at year t , denoted respectively as $r_R^d(t) = (r_{R, \mathcal{P} \cap \overline{AI}}^d(t), r_{R, \mathcal{P} \cap AI}^d(t))$ and $r_I^d(t) = (r_{I, \mathcal{P} \cap \overline{AI}}^d(t), r_{I, \mathcal{P} \cap AI}^d(t))$. In case of the a missing JSC in one of the two corpora at a given year t , we fill the missing rank by a maximum value set at 100.

With one of such couple of ranks (either for references or citations), we can locate one discipline d in a 2D space of rankings, as shown in Fig. 3. In this coordinate system, lower the value of one axis, better the associated rank, and then higher the number of collected citations along this axis. The angle θ indicates the deviation of d from the diagonal in this space, where the ranks are exactly the same in each corpus.

From this map we define the *common interest area* between $\mathcal{P} \cap AI$ and $\mathcal{P} \cap \overline{AI}$ (in grey) as the region close to the diagonal where the two rankings of the fields of study inside it are almost the same and not significantly varying over time. This zone is comprised between the two lines of respective equations $r_{\mathcal{P} \cap AI} = r_{\mathcal{P} \cap \overline{AI}} + \tau$ (above the diagonal) and $r_{\mathcal{P} \cap AI} = r_{\mathcal{P} \cap \overline{AI}} - \tau$ (below the diagonal), with τ a parameter set as 10 in this paper.

The blue zone over the common interest area is corresponding to the space where the disciplines inside it have a better rank in $\mathcal{P} \cap \overline{AI}$ than in $\mathcal{P} \cap AI$. In the case of references (citations), such disciplines are thus more cited by (citing) the former than the latter. The purple zone over the common interest area is thus the opposite case.




Fig. 3 Coordinate system to characterize the distribution of disciplines in the references and citations of the corpora $\mathcal{P} \cap \bar{AI}$ and $\mathcal{P} \cap AI$.

In Section. 4.2, we apply this per decade periods instead of single years, ie. we build decennial ranking space from the total number of citations given or received by the represented JSCs in $\mathcal{P} \cap \bar{AI}$ and $\mathcal{P} \cap AI$ within the period under study.

3.3 Capture the use of AI within the neuroscientists

3.3.1 Classification of the authors according to their AI involvement

The dataset \mathcal{P} alone would not be sufficient for a study of social dynamics of scientists inside neuroscience. Indeed the latter would have had published into disciplines or fields of research other than neuroscience or in journals that are not labeled as such in the WOS database. We thus add another level of granularity by extracting all the missing publications that have been published by the authors who belong to \mathcal{P} . Notice that these publications can have been published in the neuroscience journals extracted in our list of 421 journals (see Section 3.1), since we focus only on the more impacting authors who have published papers with at least 10 generated citations (and having at least 10 references). From this extended set of papers we deduce the duration of the scientific life of each author, and we select first those who began after 1940. In addition, because of the poor disambiguation of some authors in the MAG dump-files, whose are showing at first glance a career spanning over one or more centuries and far too many publications, we select only the authors having a reasonable duration academic life, here up to 50 years.

With these added data, we also capture the real activity (before 2019) of each author in AI research, not only in neuroscience but also in other fields of research by attributing them an AI score. We define it as the share of AI-related papers published before 2019, $f_{AI}(a) = n_a^{\mathcal{P} \cap AI} / n_a^{tot}$. To prevent the unexpected effect of accumulation of scientists having an AI score equal to 1 due to a single publication in their very short career, we filter a second time the authors dataset by considering only those who have published at least 3 papers. We thus consider in the following only 886,074 scientists, in which 188,325 (only 16% of the scientists) have published at least one AI-related paper in neuroscience. According to the f_{AI} distribution shown in Fig.4, we divide this set of authors in four parts, namely \bar{Q} ($f_{AI} = 0$; $N_a = 697749$), Q_0 ($f_{AI} \in (0, 0.5)$; $N_a = 182925$), Q_1 ($f_{AI} \in [0.5, 1)$; $N_a = 4977$) and Q_2 ($f_{AI} = 1$; $N_a = 423$).




Fig. 4 Distribution of the AI score f_{AI} of the authors.

3.3.2 Collaboration network

We then build a temporal co-signature network TCN from the whole set of papers containing both \mathcal{P} and $\overline{\mathcal{P}}$. Two authors i and j publishing together at year t are linked by an edge weighted by the number of common publications at this year $w_{ij}(t)$. All co-signatures in all papers published in year t shape a weighted snapshot called TCN(t). In order to unveil the structure of the collaboration at macro-scale in the whole period 1970-2019, we build the weighted time-aggregated co-signature network ACN, which includes the set V of all authors appearing in the dataset, and a set of weighted edges obtained by summing all the weights of the edges appearing in each snapshot of TCN, $E = \{\tilde{w}_{ij} \mid \forall (i, j) \in V^2 \mid \tilde{w}_{ij} = \sum_t w_{ij}(t)\}$. ACN is finally composed of 871,282 nodes and 7,420,423 weighted edges, and is divided into 17,599 independent connected components, with a giant one including 93% of the nodes. The smallest components exhibit also a non-negligible representation of authors belonging to the quartiles Q_1 and Q_2 , respectively 25,5% and 26,5% of the population in these quartiles. Since the authors belonging to them are supposed to drive the AI research in neuroscience, we choose to keep all the communities for our analysis, even smaller than the giant component.

Here the collaboration network is built only among the most impacting scientists of our dataset. Indeed we draw on one hand the collaboration network directly from the set of neuroscience papers that are the most impacting science in terms of citations, and on the other hand we have focused on the authors who published at least 3 articles up to 2019. Therefore we don't take into account the short-terms collaborations among very early-career scientists, small-career scientists or extra-academic ones that are involved in one or two publications with too small impact but who can potentially drive forward innovation.

3.3.3 Intensity of collaborations between f_{AI} quartiles

In this section we detail the general computation of z-scores which will be used in Section 4.4.2 to indicate the propensity of collaboration in the time-aggregated collaboration network ACN.

We first consider a static collaboration network $G = S, E$ with S the set of scientists and E the set of edges between them. Each scientist in S is identified by their respective AI score f_{AI} . According to Section 3.3.1, the distribution is divided in four

parts $\{\bar{Q}, Q_0, Q_1, Q_2\}$, from the less AI-expert to the most one. Therefore one author is belonging to one given group Q given its own AI score. By aggregating the scientists into such Q groups, we compute the 4×4 matrix τ_{obs} of effective edges between these groups, where one element is written as $\tau_{obs}(Q_i, Q_j) = |\{(u, v) | u \in Q_i, v \in Q_j\}|$.

Then we create a set of N alternative collaboration networks $\{\tilde{G}_k\} = \{(S, E_k)\}$, with $k = \{1, \dots, N\}$, which are all based on the same set of scientists but have different set of edges obtained by perturbing the network with a uniform shuffling which conserves the degree distribution of the scientists. By applying the same clustering of scientists in Q groups for each network as for the real network G , we draw a set of matrices $\{\tau_k\}$ corresponding to the share of edges between the Q groups within the randomized networks $\{\tilde{G}_k\}$. From this set we extract on one hand the average matrix τ_{sim} where one element is defined as $\tau_{sim}(Q_i, Q_j) = \frac{1}{N} \sum_k \tau_k(Q_i, Q_j)$, and on the other hand the standard deviation matrix σ where one element is defined as $\sigma(Q_i, Q_j) = \frac{1}{N-1} \sum_k (\tau_k(Q_i, Q_j) - \tau_{sim}(Q_i, Q_j))^2$.

We finally compare the empirical matrix τ_{obs} of the real network with the simulated one τ_{sim} through a z -score matrix where one element is defined as follow:

$$z(Q_i, Q_j) = \frac{\tau_{obs}(Q_i, Q_j) - \tau_{sim}(Q_i, Q_j)}{\sigma(Q_i, Q_j)}. \quad (1)$$

This standardization is used here to test the over-representation or under-representation of a given number of edges between two groups with respect to the corresponding average simulated value which is representing an ideal situation through randomization. Applying the whole procedure before on the network ACN built in Section 3.3.2, with $N = 100$, gives the Fig. 13B. Fig. 13A is also drawn by applying the same method with same N on each temporal snapshot of the temporal collaboration network TCN (see also Section 3.3.2), and by aggregating the groups Q_0, Q_1 and Q_2 into a single one called Q_i (see Fig. 13A) that includes all the scientists with at least one AI-related publication. Since TCN is undirected, the 4×4 matrices $\tau_{obs}, \tau_{sim}, \sigma$ and z are therefore reduced to scalars defined with only the two groups \bar{Q} and Q_i .

4 Results

4.1 Alignment of the neuroscience AI specialty with general neuroscience

First of all we investigate the embedding of AI technologies and knowledge within neuroscience through the dynamical interaction between the disciplinary environment of the two fields respectively.

After having divided the neuroscientific papers in two sub-corpora, one including the AI-related papers ($\mathcal{P} \cap AI$) and another including non-AI ones ($\mathcal{P} \cap \bar{AI}$), we build for each of them two disciplinary rankings which could vary over time, one for the disciplines appearing in their respective references ($r^R(t)$) and another for those appearing in their respective citations ($r^I(t)$). The reference ranking associated with one given corpus summarizes its main influential fields of study, thus its disciplinary structure on which it draws upon over time, while the citation ranking indicates the

fields that are impacted at first by the corpus. Section 3.2.1 details the building of such temporal disciplinary rankings more precisely.

We then compare macroscopically at each year the reference rankings of the two corpora ($r_{\mathcal{P} \cap AI}^R(t)$ and $r_{\mathcal{P} \cap \overline{AI}}^R(t)$) and also their citation rankings ($r_{\mathcal{P} \cap AI}^I(t)$ and $r_{\mathcal{P} \cap \overline{AI}}^I(t)$) by using a common similarity metrics J provided by Gargiulo et al. (2016, see their supplementary materials). This last measure, comprised between 0 and 1, evaluates how much similar are two given rankings r_A and r_B : if $J(r_A, r_B) = 1$, the two rankings are exactly the same, ie. containing the same elements with the same respective ranks; conversely, if $J(r_A, r_B) = 0$, the elements included in r_A are not in r_B , whatever their respective ranks.

Fig. 5 shows the evolution of two such similarity measures, a first one for comparing the rankings of disciplines cited by the corpora $\mathcal{P} \cap AI$ and $\mathcal{P} \cap \overline{AI}$ (blue curve) and a second one for comparing the rankings of the disciplines that are citing them over time (red curve).




Fig. 5 Instantaneous similarity between the references’ or impact’s rankings of the two corpora $\mathcal{P} \cap AI$ and $\mathcal{P} \cap \overline{AI}$ at year t .

Fig. 5 shows that the AI-related literature in neuroscience is impacting directly almost the same disciplines at the same level of interest shown by the non-AI corpus for all the studied time period. AI research in neuroscience is thus always being beneficial for all the same disciplines that are structuring the last field of research since the 1970’s. However the references remain quite different between both sets, but the almost linear growing of the similarity of reference is a global sign of bibliographic homogenisation as well between the two corpora. AI is thus progressively penetrating the cognitive structure of neuroscience according to this evolution.

4.2 Neuroscience and AI: two different disciplinary landscapes

In the following we compare how the individual disciplines appearing in the references and/or in the citations contribute to shape the respective disciplinary landscape of the corpora $\mathcal{P} \cap AI$ and $\mathcal{P} \cap \overline{AI}$.

By applying the framework described in Section. 3.2, we associate with each represented discipline d in the references or citations of the corpora $\mathcal{P} \cap \overline{AI}$ and $\mathcal{P} \cap AI$, their time-aggregated ranks over decades, $(\tilde{r}_{\mathcal{P} \cap \overline{AI}}^d(T), \tilde{r}_{\mathcal{P} \cap AI}^d(T))$, with $T = [t_0, t_0 + 10)$ a given time period where $t_0 \in \{1970, 1980, 1990, 2000, 2010\}$. These ranks are built with the total numbers of citations received by d (references) or given by d (citations, impact) during the time period T in the two corpora.

We represent them in the 2D space of rankings associated with each corpus respectively, as shown in Fig. 6. In one of the maps drawn on this figure, one colored disc represents a specific discipline d that is located with its respective ranks in each corpus during the considered time period, in references (Fig. 6A) or citations (Fig. 6B). Lower the value of one axis, better the associated rank, and then higher the number of collected citations along this axis. From these maps we define the common interest area between $\mathcal{P} \cap AI$ and $\mathcal{P} \cap \overline{AI}$ as the region close to the diagonal where the two rankings of the fields of study are almost the same and not significantly varying over time. These fields are colored in grey in these maps. The fields that are more dispersed around the diagonal are represented with two different colors, the blue ones having a better rank in $\mathcal{P} \cap \overline{AI}$ than in $\mathcal{P} \cap AI$ and the purple ones having a better rank in $\mathcal{P} \cap AI$ than in $\mathcal{P} \cap \overline{AI}$. All of them are abbreviated for sake of clarity. A correspondence table with their complete name in WOS are given in App. C.

To characterize each of these three area around the diagonal in the references (citations) of each corpus, we use the annual couples of ranks $r_R^d(t)$ ($r_I^d(t)$) of each represented discipline d in $\mathcal{P} \cap \overline{AI}$ and $\mathcal{P} \cap AI$ and compute the temporal angle $\theta(t)$ that indicates its temporal deviation from the diagonal, here corresponding to $\theta = 0$. Then we compute inside each zone and at each year the average angle $\langle \theta \rangle$ and the standard deviation from the latter, both represented in Fig. 7. The impact plot shows only the existence of the grey zone until the 1980's, which means the two corpora was cited by common disciplines in this period, without characteristic disciplines that are shaping further the blue and purple curves.

The disciplinary composition of the common interest area, concentrated around the diagonals of all the maps, includes rather the same fields at each decade both for reference and impact side. According to the most persistent ones over decades inside the references and generated citations of both corpora (see upper left boxes in Figs. 6A and 6B), the core of the observed citation dynamics lies in neuroscience and is composed by disciplines that are mainly associated with medicine and biomedical research, such that *Biochemistry & Molecular Biology*, *Behavioral Sciences*, *Clinical Neurology*, *Physiology*, *Cell Biology*, *Psychology*, *Psychiatry* and *Ophthalmology*. This zone is thus coinciding with the definition of neuroscience given by neuroscientists themselves, namely it “include[s] all fields that are involved with the study of the brain, the behaviors that it generates, and the mechanisms by which it does so, including cognitive neuroscience, systems neuroscience and psychology” (Hassabis et al. (2017), p. 245). This zone is also accompanied by disciplines whose ranks are more variable and that are associated to more technological aspects of neuroscience, such as *Computer Science*, *Engineering*, *Radiology*, *Neuroimaging* and *Audiology & Speech-Language pathology*. A description of the evolution of the disciplinary landscape inside this zone is detailed in App. B.

The special disciplinary ecosystem of the AI-corpus, in blue in the rankings maps of Fig.6, is also centered around biomedical fields of study that tend to be close to the disciplines that characterize the whole neuroscience as mentioned above. In particular, as shown in shown in Fig.7, the mean angles and the standard deviations of the disciplines in the associated zone are asserting a global concentration of them towards the common interest area, which includes on one hand the most influential

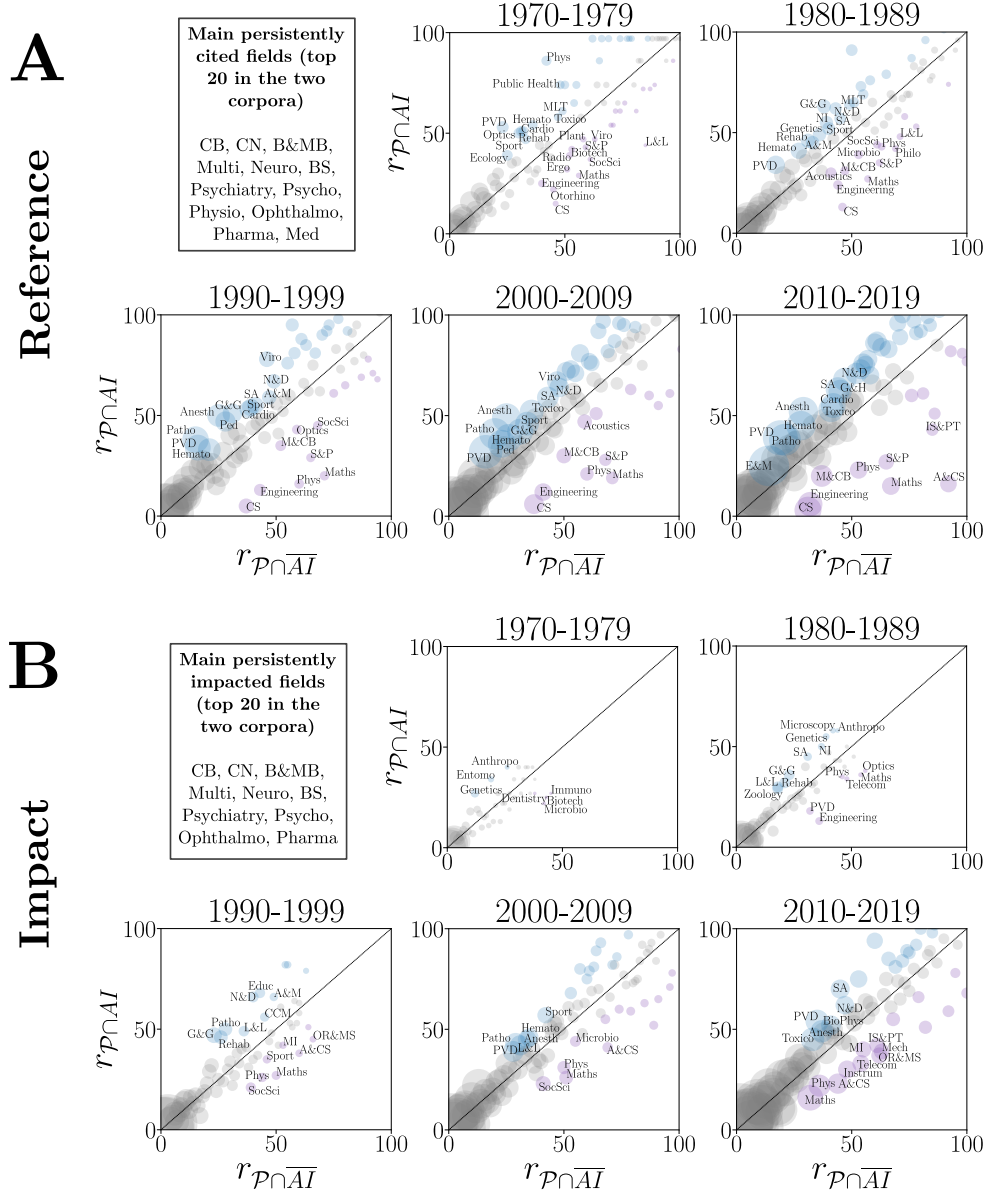


Fig. 6 Time-aggregated ranking maps of the fields of study involved in publications *cited by* both studied corpora (A) and in publications that are *citing* them (B). The dashed lines show the diagonal where the rankings are exactly the same in the corpora $\mathcal{P} \cap AI$ and $\mathcal{P} \cap \bar{AI}$. Only are shown the most significant disciplines in both corpora, with the condition that they appear in both corpora. The sizes of the discs, based on their empirical number of citations in the two different corpus, are normalized to make comparable their respective citation weights within each corpus. Grey points are corresponding to disciplines remaining confined in the area between the two lines of respective equations $r_{\mathcal{P} \cap AI} = r_{\mathcal{P} \cap \bar{AI}} + \tau$ (above the diagonal) and $r_{\mathcal{P} \cap AI} = r_{\mathcal{P} \cap \bar{AI}} - \tau$ (below the diagonal), with $\tau = 10$, and with no important variations of positions from one period to one another. The most persistent ones over decades, with a rank lower than 20, are mentioned in the upper left box of each figure A and B. Finally, blue points are the most preferred disciplines of the $\mathcal{P} \cap \bar{AI}$ corpus while purple ones are those for the corpus $\mathcal{P} \cap AI$. Abbreviations of the apparent disciplines are given in App. C.




Fig. 7 Temporal average angles produced by the disciplines in each area of reference’s (left) and impact’s (right) diagrams of Fig. 6, with respect to the diagonals represented here by a dashed line at $\langle \theta \rangle = 0$. The colored area covering the curves are representing their respective standard deviation from the mean angle. The angles are expressed in radians.

fields of research (reference) and on the other hand the core of those that show the same interest for the two studied corpora (impact). We also notice in the reference’s angle plot of Fig. 7 a deviation from the diagonal of the disciplines in this zone of references since 2010, that is also shared by the disciplines in the common interest area of references. This suggests a recent shift of references shared by the two corpora toward fields of study preferred by the non-AI corpus $\mathcal{P} \cap \overline{AI}$.

Conversely, the special disciplinary ecosystem of $\mathcal{P} \cap AI$ evolves differently by representing the mathematical, computational and technological part of neuroscience since 1970. The regular references to *Computer Science*, *Physics*, *Statistics & Probability*, *Mathematical & Computational Biology* and *Engineering* show a large influence of technological-oriented research in this particular AI research in neuroscience. Progressively concentrated towards the common interest area between 1970 and the late 1980 (see Fig. 7), these references preferred by $\mathcal{P} \cap AI$ become more further away from the common interest area and more dispersed after this period, thus indicating a cognitive differentiation of references on which the AI-related corpus $\mathcal{P} \cap AI$ is drawing upon from the non-AI corpus $\mathcal{P} \cap \overline{AI}$. In addition, while neuroscience and associated medical fields – as *Clinical Neurology* and *Neuroimaging* – remain the primary stakeholders in the AI research conducted within it, the latter appears to be of varied interest to a subset of disciplines since 1980’ that do not place as much emphasis on works in the non-AI corpus, such as *Mathematics*, *Physics*, *Computer Science* and also *Automation & Control Systems* (see Fig. 6B). According to Fig. 7, these citing disciplines are also closer to the common interest area than in references, thus confirming the involvement of AI research into the general scientific purposes of neuroscience.

By regarding the disciplines with the most significant evolution within the AI ecosystem, Fig. 8 shows the rise of *Neuroimaging* and *Radiology* both in references and impact sides, as already observed for the common interest area, and also the progressive domination of *Computer Science*. This figure shows also the decrease of the influence of fields of research linked to the disciplinary lines of the neuroscience and of the corpus $\mathcal{P} \cap \overline{AI}$. This is also a sign of a progressive differentiation of the research supported by the AI research from the rest of neuroscience. The case of physiology, which contributed to the foundations of neuroscience (Cooper & Shallice, 2010), is particularly eloquent with its distancing in time from the highest positions in the ranking since the 2000’s. This observed retreat of master neuroscience disciplines, that had a strong influence in the building of AI’s in this field since the 1970’s, suggests that

a social transformation occurred in this special research during the studied temporal period. Perconti and Plebe (2020) mention such a transformation, in which AI in neuroscience was a matter for biomedical specialists before becoming an object of study and technological developments for engineers. This will be shown empirically in the following sections.

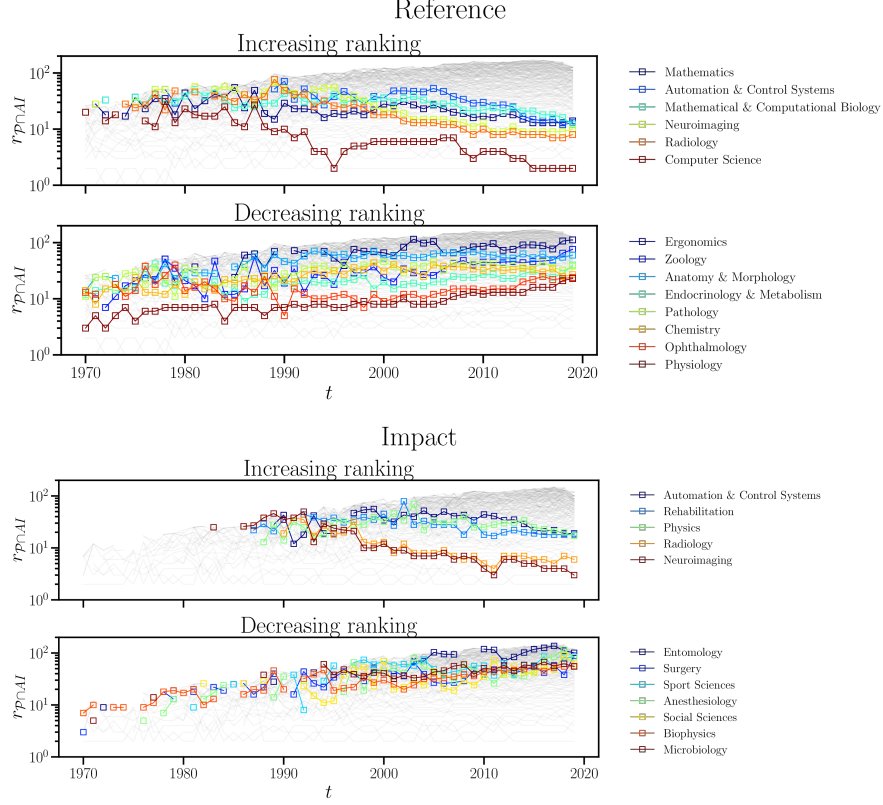


Fig. 8 Time evolution of the ranking $r_{\mathcal{P} \cap AI}$ of the disciplines that are mostly represented by the corpus $\mathcal{P} \cap AI$ in the references (top) and impact sides (bottom), ie. situated in the lower zone of the maps shown in Fig.6. Only the most significant are shown with colors.

4.3 An AI literature confined in a small set of journals

Another factor of differentiation of the AI research community from the core of neuroscience is the set of journals in which the first are mainly published. Starting on the set of 421 journals included in the whole neuroscience dataset, we compute the temporal share $a_{AI}(t)$ of AI-related publications in each of them (called also *temporal AI activity*), and its global share a_{AI} of such publications since its launch year (called also *justAI activity*). Here we consider the whole publication history of each journal for computing these scores, since we omit some papers published in them with the

filter based on a minimum number of references and impact (see Section. 3.1). These scores are represented in Fig. 9. We also extract the journals with a_{AI} higher than 10%, whose names and launch years are given in the table in the bottom of this figure.


According to the evolution of the AI activities of neuroscience journals, the AI-related research in this field is concentrated around a small subset of journals providing development in computational techniques mainly linked to neural networks and cognition. These journals are quite representative of the connectionist wave of AI which seems to be active in neuroscience since the late 1980's (Cardon et al., 2018). They are representing 32,7% of the scientific production of AI in neuroscience during the whole period 1970-2019, according to our dataset, the rest being distributed among the 398 other journals of the dataset with a much smaller AI-activity (less than 10%).

The spectrum of a_{AI} (top yellow curve on Fig. 9) shows that the launches of the most AI active journals in the neuroscience field are concentrated around three periods. The first one spans the period 1987-1994 with especially 7 journals (no. 3 to 9 in the table of Fig. 9) whose a_{AI} is higher than 28% and whose scopes are oriented toward computational neuroscience and the use of neural network formalism for complex calculations. These journals, except no. 7 and 9 (respectively *Network: Computation in Neural Systems* and *Neural Network World*) are showing well-sustained activity on AI research until 2019. This period falls commonly into the second "AI winter", when AI-research funding and the production of scientific results and associated technological solutions were at their lowest for a second time (Cardon et al., 2018; Schuchmann, 2019). Paradoxically, neuroscience are especially active in such research in this period, as shown by the strong, long-lasting AI-activity of these journals created in this period.

The second period includes the journals 10, 11 and 12 that have been launched between 1999 and 2001. The journal *IEEE Transactions on Neural Systems and Rehabilitation Engineering* is especially oriented toward the development of computational methods and technological tools to capture the neural activity of the brain. This period is followed by another until 2007 which did not lead to the launch of AI-active journals.

Finally the third period following the latter, spanning the period 2007-2009, is composed of 7 journals (no. 14 to 20) that are fostering research at the crossroads of neuroscience and cognitive science. It includes especially the most AI active journal of our dataset, *Evolutionary Intelligence*, whose scope is oriented toward evolutionary computation, which is a subset of the field of optimization. Its temporal activity is also strongly sustained until 2019.

Aggregating the date over time, Fig.10 shows a linear correlation between the mean AI-activity in the journals and the average of the AI scores f_{AI} of the authors that have published at least one paper in these journals. The tail at highest a_{AI} (higher than 0,1) is corresponding to the top 15 of the most AI-active journals given in Fig. 9. This particular result thus unveils an attraction of the authors publishing the most AI-related works of our dataset in the journals with a high AI-activity in time. This suggests that AI research in neuroscience is done by a specialized scientific community inside this field with its own journals for communicating results.



Journal	Launch year in the database	a_{AI}
1 Cognitive Science	1976	0.116029
2 Journal of Sensory Studies	1986	0.104065
3 Neural Networks	1987	0.530183
4 Neural Processing Letters	1988	0.694971
5 Neurocomputing	1989	0.559622
6 Neural Computation	1989	0.370334
7 Network: Computation in Neural Systems	1990	0.281250
8 Adaptive Behavior	1992	0.313070
9 Neural Network World	1994	0.392226
10 Cognitive Systems Research	1999	0.303797
11 IEEE Transactions on Neural Systems and Rehabilitation Engineering	2001	0.224457
12 Neuroinformatics	2001	0.184906
13 Plos Computational Biology	2005	0.122665
14 Frontiers in Neuroinformatics	2007	0.185841
15 Computational Intelligence and Neuroscience	2007	0.547596
16 Cognitive Neurodynamics	2007	0.242938
17 Frontiers in Computational Neuroscience	2007	0.226304
18 Evolutionary Intelligence	2008	0.695946
19 Cognitive Computation	2009	0.513138
20 Topics in Cognitive Science	2009	0.100173
21 Journal of Mathematical Neuroscience	2011	0.126984
22 NeuroImage: Clinical	2012	0.111653
23 Biological Psychiatry: Cognitive Neuroscience and Neuroimaging	2016	0.112500

Fig. 9 Top diagrams: A: Temporal AI activity of the journals. One square at the position (j, t) of the diagram is the share of AI publications in the journal j at time t . One column of the plot is the AI activity of one journal. B: Distribution of the global AI activity of the journals over all their publications up to 2019. The horizontal axis is corresponding to the journals axis of the bottom plot A, as shown by the red dashed zone. C: Zoomed activities of the 23 most active journals in AI research, with a global AI-activity higher than 10%. They are indicated in the bottom table.




Fig. 10 Correlation plot between the global AI activity a_{AI} of the journals appearing in \mathcal{P} and the average AI scores \bar{f}_{AI} of the authors that have published in it between 1970 and 2019. Each point represents one journal in the horizontal a_{AI} space, colored by the standard deviation of the f_{AI} scores of the authors who have published in it. The maximum average value of \bar{f}_{AI} is situated around 0,4.

4.4 Who are the scientists making AI in neuroscience?

4.4.1 Disciplinary profiles in each quartile of the f_{AI} distribution

In this part we focus on the profiles of the scientists doing AI in neuroscience (in Fig.4, included in one of the quartiles Q_0 , Q_1 or Q_2) and how they are inserted into the global authorship landscape of neuroscience. In particular, we compare these last with the scientists in neuroscience who did never publish AI-related papers (included in the quartile \bar{Q}) under two aspects, namely (1) their disciplinary background and (2) their disciplinary trajectory in academia.

We define the “disciplinary background” of one author the combination of disciplines corresponding to the papers he/she has published in his/her first year of academic life. Then we compute for each quartile the temporal cumulative number of new scientists formed in each represented disciplinary background in this quartile, as shown in Fig.11. In this way we assess for each quartile the main native specialties in which our authors have first published.

Fig.11 shows that the profiles included in Q_0 and \bar{Q} overall the period 1940-2019 are very similar. They are mainly confined in biomedical research around neuroscience, as well as in *Multidisciplinary Sciences* which is represented by 91 international journals with a broad topical diversity. We recognize also the main fields of research that are shaping the common interest area of the disciplinary ecosystems of the two separated AI and non-AI corpora, that are shown in the grey rows of the tables in Fig.6. These two profiles are thus representing at most the disciplinary spectrum of neuroscience itself. We notice in addition the spectacular increase of the *Neuroscience* curve in the two plots until the 1970’s (after having emerged in 1957 for \bar{Q} and in 1962 for Q_0), that is followed by a quasi-linear progression until 2015. This boom of neuroscience profiles in these quartiles let us suppose that the modern neuroscience is progressively inserting institutionally as a well-structured discipline in science.

The profiles of Q_1 and Q_2 shown in Fig.11 are at the opposite of the previous ones, coming at most from fields of research related to *Computer Science*. The specialty *Computer Science*, *Neuroscience*, which emerged in 1988 for Q_1 and in 1991 for Q_2 indicates the rising of a group of scientists that are specialized into computation in neuroscience. Notice that the behavior of the *Neurosciences* curve within

Q_1 Fig.11 appears later than computational and engineering profiles, suggesting also that AI-related knowledge and technological tools penetrate progressively the global neuroscience field.

From the subset of authors belonging to the most frequent disciplinary background in each quartile, we consider their respective disciplinary profile as the disciplines corresponding to the journals in which they publish throughout their scientific life – that would not be ended for the youngest still publishing in 2019. We therefore draw from these two parameters – disciplinary background and career-related disciplines – the typical disciplinary trajectories in each quartile, which are shown in Fig.12. Since the backgrounds are built from the publications at a given year, the authors who began their career in 2019 would have a disciplinary profile corresponding to it. For avoiding an over-representation of some confined disciplinary trajectories due to these newcomers, we therefore consider only the authors who began at most in 2018. For the sake of clarity we select only the most significant trajectories.

As a confirmation of the disciplinary background shown in Fig.11, Fig.12 shows that the studied authors in \overline{Q} and Q_0 are involved in a similar disciplinary landscape centered around *Ophthalmology*, *Clinical Neurology* and general *Neurosciences*, while Q_1 and Q_2 are more confined into *Computer Science* and *Engineering*. In particular, by considering again neuroscience as the subset of disciplines including *Neurosciences*, *Clinical Neurology* and *Neuroimaging*, we observe that \overline{Q} and Q_0 are more involved in that field of research, with respectively 77% of the authors in the first and 78% of those in the second who are showing a disciplinary profile that includes one or more of the fields of research associated with neuroscience. On the contrary, 45% of the authors in Q_1 and 42% of the authors in Q_2 have a disciplinary profile that includes such JSCs. These two quartiles are therefore mainly detached of the neuroscience goals given the disciplinary backgrounds and profiles of their main respective authors. These results show at first glance that neuroscience is a field that bring together heterogeneous profiles who seem to serve different epistemic objectives inside and outside neuroscience (Sedooka, Steffen, Paulsen, & Darbellay, 2015).

Nonetheless we notice two special things about these last quartiles. First the authors who began into the specialties including one or more neuroscience-related disciplines tends to continue in the same field of research, which is sometimes interdisciplinary such that *Computer Science*, *Neurosciences* and *Neuroimaging*, *Neurosciences*, *Radiology*. This means that these scientists are formed into a disciplinary context centered around neuroscience. Second, some authors who began into *Computer Science*, *Engineering* or mathematics-related disciplines – who seem to be not concerned about neuroscience in their disciplinary backgrounds – continue into neuroscience, as shown by the combination of their originating JSCs with neuroscience ones in their disciplinary profiles in Fig.12, for instance *Computer Science*, *Engineering* leading to *Computer Science*, *Engineering*, *Neurosciences*. The late emergence of interdisciplinary profiles in computer science and neuroscience, who are more involved in neuroscience in general such as Q_2 , also testify that AI community is taking root in the global neuroscience landscape, the more recent profiles in neuroscience becoming

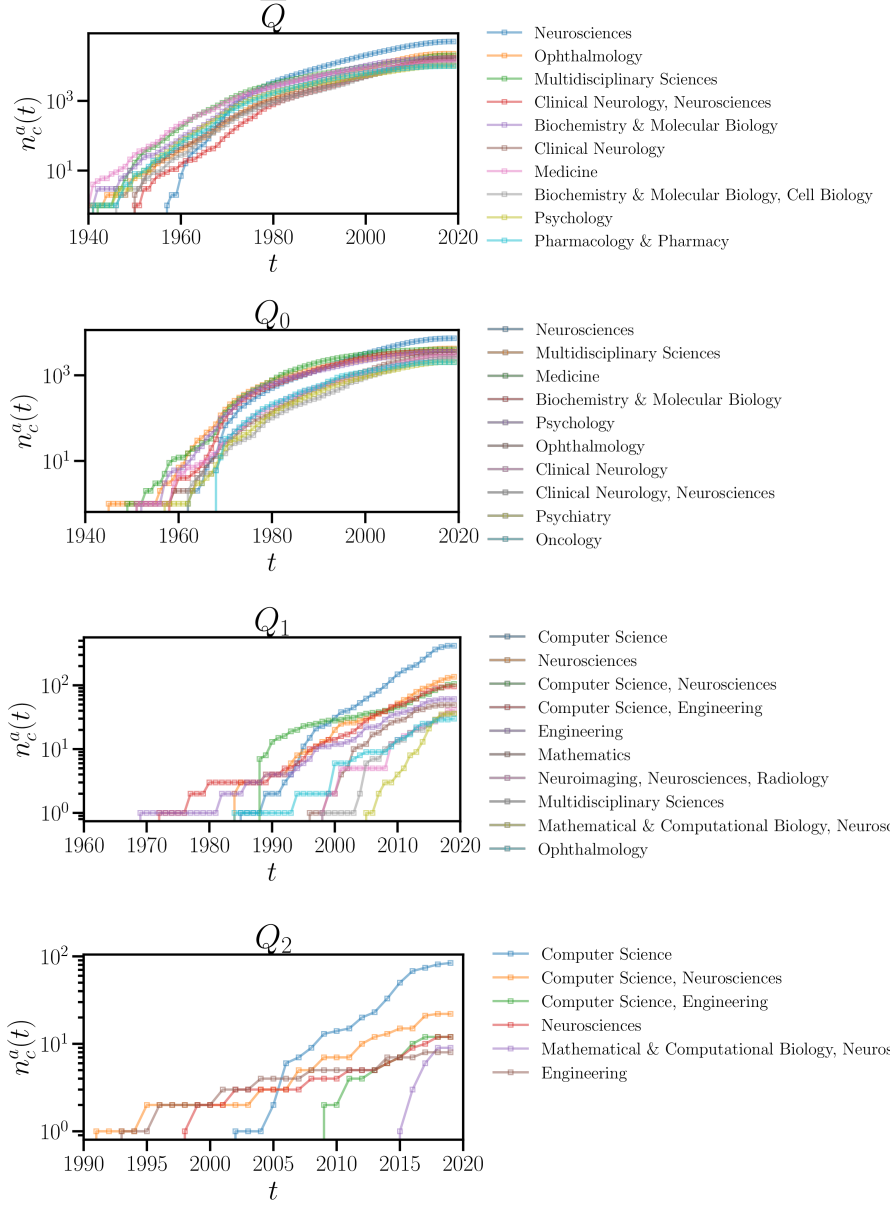


Fig. 11 Cumulative number of new authors per disciplinary background for each quartile. For one plot, the point of the curve of the specialty d at year t is the number of authors who have published their first articles inside d during their very first year of academic career, since the year of apparition of d . Only the top 10 native specialties in 2019 are shown in \bar{Q} , Q_0 and Q_1 , and the top 6 for Q_2 because of the insignificance of the following ones.



Fig. 12 Significant disciplinary trajectories of the authors confined in each quartile. On each plot lie on the left the most significant disciplinary backgrounds, and on the right lie the career-related JSCs up to 2019.

insiders in this new technological specialty. All these scientists thus represents a special labor force for neuroscience whose main expertise lies in AI, and more generally in mathematical, computational and technological tools (Perconti & Plebe, 2020).

4.4.2 Structure of the collaboration network between quartiles

Now we investigate how the AI practitioners are distributed among the neuroscience community, and how are shaped the collaborations between them, especially between the different kind of scientific profiles described before.

We first consider the temporal collaboration network TCN built in Section 3.3.2, and we evaluate the temporal standardized share of edges between scientists belonging respectively to \bar{Q} and to all other quartiles Q_i 's (see Section 3.3.3 for the calculation of such score). This score is showed in Fig.13A, from which we deduce two facts.

First, all its values are lower than 0 over time, meaning that the links between the scientists in the two studied ensembles of authors in the period 1970-2019 are under-represented compared to situations obtained by randomized reshuffles of the edges between them. Second this situation becomes worse over time, with a continuous reduction under 0 until the 1990's followed by a quick one until 2015. Although this tendency is reversed since 2015, these results indicates that the neuroscientists making AI, ie. in Q_i 's, are shaping an almost independent community inside neuroscience by widening progressively a gap with other neuroscientists belonging to the ensemble \bar{Q} . According to the profiles shown in Fig.11, since the scientists who published AI-related research until the 1990's are mostly coming from Q_0 , and since they are representing the whole neuroscience field, the result advanced above let us suppose that these scientists in Q_0 are *outsiders* in the history of AI in neuroscience, by analogy with the notion first advanced by Crane (1969) in the sociology of science.

This social divide is confirmed with Fig.13B, which represents the panorama of the links shared between the different quartiles in ACN (see also Section 3.3.3). This diagram represents a stabilized situation observed since the end of the 1990's (see Fig.D3 in App.D), where the Q_i 's cosign more together than with \bar{Q} while the scientists in the second ensemble prefer to collaborate together as well. However, in regards with the temporal similarity index of impacted JSCs shown in Fig.5, as well as the temporal evolution of the disciplinary landscapes shown in Fig.6, the Q_i 's export their knowledge mainly in the whole neuroscience, even those cited by articles written by scientists belonging to \bar{Q} . We thus assist to a social separation in neuroscience around AI research, that does not produce a strict knowledge divide in neuroscience though.

We also observe in Fig.13B a polarization within the subset of AI practitioners in neuroscience. Indeed the authors belonging to Q_1 and Q_2 are more strongly connected together than with those in Q_0 . This can be explained by the disciplinary proximity of the authors in Q_1 and Q_2 observed at first glance in Fig.11, especially the prominence between the late 1980's and the early 2000's of full computer scientists and hybrid profile publishing in journals labeled as *Computer Science*, *Neurosciences*. Furthermore, the links between Q_1 and Q_0 are much more important than those between Q_2 and Q_0 . The scientists in Q_1 thus appears to be the most interdisciplinary by assuring the bridge between these differentiated profiles inside AI community. These particular

outsiders described before are driving the diffusion of AI in neuroscience from computation to medical and clinical applications, given their disciplinary trajectories shown in Fig.12.

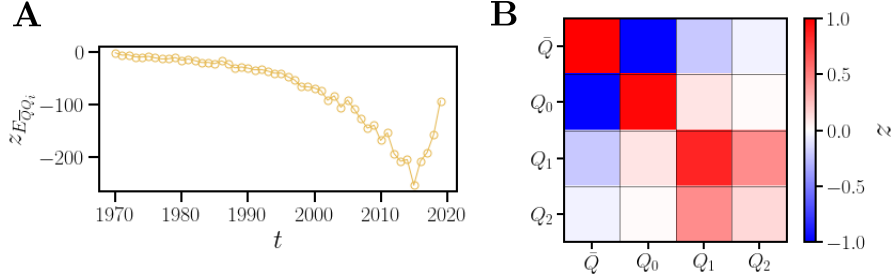


Fig. 13 A: Temporal z -score applied to the edges in TCN between the scientists in \bar{Q} and those in the other Q_i 's, here aggregated together under the notation Q_i . B: Z -score matrix of the edges between each quartile in the time-aggregated collaboration network ACN. The values are normalized with the absolute maximum one in the matrix.

5 Discussion

In this article we have explored the penetration of AI in neuroscience through the co-development of both fields and especially the construction of a technological specialty centered around AI inside neuroscience. To do this we conducted a scientometric analysis of an exhaustive bibliometric database that aims to represent at best the neuroscience research between 1970 and 2019. This analysis lies on many indicators.

First, with a comparative analysis of the egocentric citation networks associated with respectively AI research and non-AI one, we have shown a progressive cognitive differentiation of the first from the second, which is based on the employed bibliographic references in the two fields respectively. We observed indeed a specialization of AI research toward computer science, mathematics and engineering, while the core of neuroscience research draws upon biomedical and clinical research fields related to it. This differentiation is also shown by the neuroscience journals landscape, in which a small set of 23 journals, the most active in AI research in neuroscience, are representing those STEM fields in references and are gathering the scientists who are the most active in AI in the neuroscience field as well. However, we showed that the demarcation of AI research that is observed in bibliographical references remains partial because its citational impact is more broadly distributed across the entirety of the field of neuroscience. The impact of AI-related research provided by neuroscience is thus not confined into a subset of research which could be the same as the one visible in its employed references – even if AI research in neuroscience exhibits a particular accent for engineering and medical imaging technologies in its references and target fields of research since around the 1990's, just after a period marked by a decline in AI research activities called “second AI winter”.

In a second step we explained this partial specialization by conducting a relational approach based on the scientific profiles of the authors of the studied database.

By distinguishing the AI practitioners in neuroscience from the neuroscientists, we have shown that the first tend to not maintain links with the second in the temporal co-authorship network which includes the main collaborations since 1970. With the previous results we concluded that AI researchers are “outsiders” in the field of neuroscience. We especially have distinguished two classes in this group of scientists, namely a first wave of outsiders including authors who are trained in the main disciplines that are shaping the neuroscience field since the 1940’s and who have a low AI activity, and a smallest second wave which emerged around the 1980’s and that is including authors who are trained in other disciplines that are not represented in the former group, such as computer science and engineering, and who are exhibiting the highest activities in AI research in general (inside and outside neuroscience). Furthermore this second group is not the most involved in the field of neuroscience, its members keeping to publish within their original disciplines that are mainly STEM ones.

This social polarization inside the AI practitioners in neuroscience suggests that AI becomes over the years a set of technologies that need to be shaped not only by neuroscientists themselves, but also with the help of scientists who come from outside neuroscience, or from within but with an interdisciplinary background, and who present specific expertise about AI itself. Since we considered AI as a global research-technology in science (Hentschel, 2015; Marcovich & Shinn, 2020; Shinn & Joerges, 2002), these results thus are a consequence of the diffusion of AI outside its originating STEM disciplines and throughout the science system (Gargiulo et al., 2022). The second wave of AI researchers described before is indeed representing the mobility of experts towards other fields of research in order to propose and integrate the associated knowledge and technologies to achieve some of their disciplinary objectives, or, less ambitiously, solve some technical problems that could not be solved within the intrinsic epistemic framework of the welcoming discipline.

These results thus illustrate quite well the generic character of AI when applied in neuroscience, which is producing a social and cognitive differentiation inside the latter (Shinn & Joerges, 2002). However we could ask the relevance of this criterion in this case study, where only 3% of the neuroscience publications in our dataset is concerned about AI according to our keyword filter introduced in Section 3.1. This leads especially to ask how spread is AI across the topics covered by neuroscience over time, ie. if AI is present in all subfields of neuroscience or concentrated around a few ones. With such a topic space, and by reusing the data about the authors and the co-signature network, we could evaluate more precisely how much universal is AI inside neuroscience through the distribution of authors in this topic space, and then deduce the propensity of AI to fit with the knowledge and methods associated with some topics rather than others.

Nonetheless, the dynamical process of integration of AI in neuroscience is exhibiting some differences regarding the global history of AI in science depicted in (Gargiulo et al., 2022). Although the development of AI in science is statistically characterized by a disciplinary closure around STEM disciplines between 1980 and 2010, AI was keeping to be invested and developed by neuroscience for its own epistemic purposes in this period, even marginally, as shown by a large number of journals created in this period and that are very active in AI research (see Section 4.3), and also its citational impact

in general (see Sections 4.2 and 4.1). However, a social closure was simultaneously occurring and accentuated inside the field, where not everyone is finally using AI at all, even in recent days. According to Shinn and Joerges (2002), this is typical of a differentiation process observed in the conception phase of an instrument.

The penetration of AI inside a single discipline thus could also be described as another underlying dynamical process of development of the associated knowledge and instruments following the four steps of Hentschel (2015), which is included inside the diffusion phase of the original instrument. For instance, the developed AI inside neuroscience, which could be slightly different from the originally produced one in STEM disciplines, would be also adapted to furnish other capabilities in neuroscience first, but also in other disciplines of fields of research afterwards if they judge it useful for their own goals. Under such an hypothesis, we could explore more precisely one one hand the expansion of adjacent possible of neuroscience caused directly by AI, and on the other hand the expansion of the adjacent possible of other fields that have been influenced later by an adapted AI that was designed in neuroscience (Bianchini et al., 2022; Kauffman, 2000; Monechi et al., 2017).

In addition of the limitations mentioned throughout this paper, especially for the building of our database described in Section 3, these results are facing another important one, which is lying into the type of publications that could differ from one discipline to another, and the manner it impacts different disciplinary community as well. Indeed we have shown that most of AI research published in peer-reviewed journals in neuroscience is impacting mainly the latter itself, but we could ask this pattern if we add other media for communicating research results, such as conference proceedings or preprints sharing platforms such as *arXiv*, which are most distributed among mathematicians, physicists and also computer scientists (Wainer et al., 2013), who are here the most active scientists in AI research. These media thus are susceptible to be chosen by scientists who are also the most active in AI research in neuroscience, as we shown in all this paper. Because of the remaining involvement of these scientists into their disciplines of training, which are mainly not neuroscience (see Section 4.4), this new consideration could lead to a further study about the impact of the AI research done in neuroscience, especially if it is exported outside this field with other communication media. This study would also brings issues about interdisciplinarity itself through the communication media, as well as the potential conflict of publishing traditions in neuroscience that could lead to an impact of these AI research toward specific disciplines outside the neuroscience purposes.

Finally, in order to study broadly global construction of AI in all the science, this paper intends to be a road map for further studies of the diffusion of AI in other disciplines or fields of research that are receptive to it, but probably with different patterns. A comparative work would then be required.

6 Declarations

- Ethics approval and consent to participate: Not applicable.
- Consent for publication: Not applicable.

- Availability of data, materials and codes: The datasets used during the current study, and the codes that permitted to analyze them, are available from the corresponding author on reasonable request.
- Competing interests: The authors declare that they have no competing interests.
- Funding: S.F. is funded through a CNRS-MITI PhD grant within the project “Epistemic Impact of Artificial Intelligence in Science” (EpiAI). This research has been partially supported by the ANR grant ScientIA (ANR-21-CE38-0020).
- Authors’ contributions: S.F. collected and analyzed the data, discussed the results, wrote the manuscript. F.G. conceived the research, collected the data, discussed the results, wrote the manuscript. M.D. discussed the results, wrote the manuscript. P.T. discussed the results, wrote the manuscript. All authors have read and approved the manuscript.
- Acknowledgements: Not applicable.

Appendix A Self-similarities of references and citations, and disciplinary concentration in the AI and non-AI corpora

This section focuses on the cognitive development of the research field associated with the AI corpus ($\mathcal{P} \cap AI$) and the non-AI one ($\mathcal{P} \cap \overline{AI}$) corpora, based on the same similarity measure introduced in Section 4.1. We compute for each corpus two such measures that could vary over time:

1. the temporal self-similarity of one given ranking r , namely the similarity between this ranking at time t and the the same at the previous time $t - 1$, denoted as $J(r_{t-1}, r_t)$; we apply it for references’ and citations’ rankings in each corpus,
2. the temporal disciplinary concentration of one corpus, namely the similarity between its references ranking r_R and its generated citations one r_I at a given time t , denoted as $J(r_R(t), r_I(t))$,

The evolution of these two different indices are represented in Fig. A1A and B respectively.

According to Fig. A1A, the references on which the AI corpus and the non-AI one draw upon respectively (solid lines) are both consolidating with time toward their highest respective values in 2017, but not at the same speed. Indeed the non-AI corpus lies on almost the same set of disciplines from year to year, while the AI-related one grow from a lower similarity in 1970 around 50% to a stable set with a similarity around 90% between 2015 and 2017. The observed decrease after 2017 is mainly due to the lack of data grabbed by MAG. The self-similarities of citations in each corpus (dashed lines) follow the same trend, but with more variations until a stable phase since 2009 and with a lower similarity level than for references. This means that the short-term impacted fields in the two corpora are also consolidating toward rather the same ones in recent days. More precisely, Fig. A1B indicates the growing concentration of references’ and impact’s rankings inside each corpus since around 1990,

ie. the impacted disciplines and those appearing in references become more and more similar, hence a research inspired by itself toward itself. This shows also the progressive disciplinary homogenisation of the two researches independently, and therefore a consolidation of the entire research field of neuroscience.

We notice also in Fig. A1 that the self-similarity of citations of the two corpora under study are almost identical, thus confirming the high constant similarity of citations' rankings between the corpora shown in Fig. 5 in Section 4.1.




Fig. A1 A: Temporal self-similarity of disciplinary rankings associated with either the references (R) or generated citations (I) of one specific corpus, either the AI-related one ($P \cap AI$) or the non-AI one ($P \cap \overline{AI}$). More precisely, one point at time t is the similarity between the ranking at time t and the former one at $t - 1$. B: Instantaneous similarity between the references used by one corpus and its produced citational impact.

Appendix B Disciplinary composition of the common interest area of the AI-related and non-AI corpora

This section focuses more precisely on the disciplines inside the common interest area of the AI and non-AI corpora defined in Section 4.2). According to Fig. 6, each discipline d appearing in the references and/or in the citations of these corpora is located in a 2D space by the coordinates $(r_{P \cap AI}^d, r_{P \cap \overline{AI}}^d)$ associated with its respective ranks in the AI and non-AI corpora. From these coordinates we compute the distance of the disciplines from the origin of the map (point with ranks (0,0)), denoted as ρ in the following. Fig.B2 shows the evolution of this distance for some disciplines, here those that are exhibiting the most significant variations (increasing or decreasing) over the years since the 1970's.

This figure shows especially a spectacular rise since around 1995 of the references' rankings associated with *Radiology* and *Neuroimaging* in both AI and non-AI corpora, as well as a growing impact of neuroscience articles published in this period on these disciplines. The 1980's and early 1990's marked also the launching of the first body-scanner machines applying positron emission tomography (PET) and functional magnetic resonance imaging (fMRI), that both enabled important discoveries on functional biological mechanisms in human brain that are induced by complex cognitive tasks (Cooper & Shallice, 2010). This trend thus testifies to the diffusion of these technologies in scientific and medical practices associated with neuroscience.

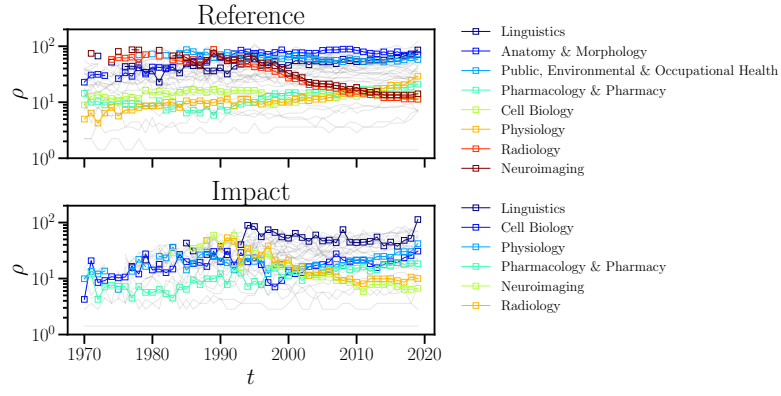


Fig. B2 Time evolution of the distance ρ of the disciplines included in the common interest area defined and shown in Fig. 6. Only the most significant are shown with colors, here with the criterion that they distance must vary at least with an increase or decrease equal to 15 between the beginning and the end of their respective trajectory.

Appendix C Web of Science categories' abbreviations table

WOS Categories	Abbreviations
Acoustics	Acoustics
Anatomy & Morphology	A&M
Anesthesiology	Anesth
Anthropology	Anthropo
Automation & Control Systems	A&CS
Biophysics	BioPhys
Biotechnology & Applied Microbiology	Biotech
Cardiac & Cardiovascular System	Cardio
Computer Science	CS
Critical Care Medicine	CCM
Dentistry	Dentistry
Ecology	Ecology
Education	Educ
Endocrinology & Metabolism	E&M
Engineering	Engineering
Entomology	Entomo
Ergonomics	Ergo
Gastroenterology & Hepatology	G&H
Genetics & Heredity	Genetics
Geriatrics & Gerontology	G&G
Hematology	Hemato
Imaging Science & Photographic Technology	IS&PT
Immunology	Immuno
Instruments & Instrumentation	Instrum
Language & Linguistics	L&L
Mathematical & Computational Biology	M&CB
Mathematics	Maths
Mechanics	Mech
Medical Informatics	MI
Medical Laboratory Technology	MLT
Microbiology	Microbio
Microscopy	Microscopy
Neuroimaging	NI
Nutrition & Dietetics	N&D
Operations Research & Management Science	OR&MS
Optics	Optics
Otorhinolaryngology	Otorhino
Pathology	Patho
Pediatrics	Ped
Peripheral Vascular Diseases	PVD
Philosophy	Philo
Physics	Phys
Plant Sciences	Plant
Public, Environmental & Occupational Health	Public Health
Radiology	Radio
Rehabilitation	Rehab
Social Sciences	SocSci
Sport Sciences	Sport
Statistics & Probability	S&P
Substance Abuse	SA
Telecommunications	Telecom
Toxicology	Toxico
Virology	Viro
Zoology	Zoology

Appendix D Temporal evolution of the inter-quartile collaboration network

References

- Andler, D. (1990). Connexionnisme et cognition: À la recherche des bonnes questions. *Revue de Synthèse*, 111(1-2), 95–127,
- Arencibia-Jorge, R., Vega-Almeida, R.L., Jiménez Andrade, J.L., Carrillo-Calvet, H. (2022). Evolution and Multidisciplinarity of Artificial Intelligence Research. *Scientometrics*, 127(2), 1–20,
- Baruffaldi, S., van Beuzekom, B., Dernis, H., Harhoff, D., Rao, N., Rosenfeld, D., Squicciarini, M. (2020). *Identifying and measuring developments in artificial intelligence: Making the impossible possible*. (OECD Science, Technology and Industry Working Papers)
- Bianchini, S., Müller, M., Pelletier, P. (2022). Artificial intelligence in science: An emerging general method of invention. *Research Policy*, 51(10), ,
- Cardon, D., Cointet, J.-P., Mazières, A. (2018). Neurons spike back. The invention of inductive machines and the artificial intelligence controversy (L. Carey-Libbrecht, Trans.). *Réseaux*, 211(5), 173–220,
- Cooper, R.P., & Shallice, T. (2010). Cognitive Neuroscience: The Troubled Marriage of Cognitive Science and Neuroscience. *Topics in Cognitive Science*, 2(3), 398–406,
- Crane, D. (1969). Social Structure in a Group of Scientists: A Test of the "Invisible College" Hypothesis. *American Sociological Review*, 34(3), 335–352,
- Frank, M.R., Wang, D., Cebrian, M., Rahwan, I. (2019). The evolution of citation graphs in artificial intelligence research. *Nature Machine Intelligence*, 1(2), 79–85,
- Färber, M. (2019). The Microsoft Academic Knowledge Graph: A Linked Data Source with 8 Billion Triples of Scholarly Data. C. Ghidini et al. (Eds.), *The Semantic Web – ISWC 2019* (Vol. 11779, pp. 113–129). Springer International Publishing.
- Färber, M. (2020). *Microsoft Academic Graph records*. (Available at <https://zenodo.org/record/3936556>)

- Gao, J., & Wang, D. (2023). *Quantifying the benefit of artificial intelligence for scientific research*. (Preprint at <https://arxiv.org/abs/2304.10578>)
- Gargiulo, F., Caen, A., Lambiotte, R., Carletti, T. (2016). The classical origin of modern mathematics. *EPJ Data Science*, 5(1), 26,
- Gargiulo, F., Fontaine, S., Dubois, M., Tubaro, P. (2022). *A meso-scale cartography of the AI ecosystem*. (Preprint at <http://arxiv.org/abs/2212.12263>)
- Gopinath, N. (2023). Artificial intelligence and neuroscience: An update on fascinating relationships. *Process Biochemistry*, 125, 113–120,
- Haas, P.M. (1992). Introduction: Epistemic Communities and International Policy Coordination. *International Organization*, 46(1), 1–35,
- Hassabis, D., Kumaran, D., Summerfield, C., Botvinick, M. (2017). Neuroscience-Inspired Artificial Intelligence. *Neuron*, 95(2), 245–258,
- Hentschel, K. (2015). A periodization of research technologies and of the emergency of genericity. *Studies in History and Philosophy of Modern Physics*, 52, 223–233,
- Kauffman, S.A. (2000). *Investigations*. Oxford University Press.
- Lake, B.M., Ullman, T.D., Tenenbaum, J.B., Gershman, S.J. (2017). Building Machines That Learn and Think Like People. *Behavioral and Brain Sciences*, 40(e253), 1–72,
- Liu, N., Shapira, P., Yue, X. (2021). Tracking developments in artificial intelligence research: constructing and applying a new search strategy. *Scientometrics*, 126(4), 3153–3192,
- Marcovich, A., & Shinn, T. (2020). Science research regimes as architectures of knowledge in context: A ‘longue durée’ comparative historical sociology of structures and dynamics in science. *Social Science Information*, 59(2), 310–328,
- McCarthy, J. (1981). Epistemological Problems of Artificial Intelligence. *Readings in Artificial Intelligence* (pp. 459–465). Elsevier.

- Monechi, B., Ruiz-Serrano, Ã., Tria, F., Loreto, V. (2017). Waves of novelties in the expansion into the adjacent possible. *Plos One*, 12(6), ,
- Perconti, P., & Plebe, A. (2020). Deep learning and cognitive science. *Cognition*, 203, 12,
- Roth, C. (2008). Réseaux épistémiques : formaliser la cognition distribuée. *Sociologie du Travail*, 50(3), 353–371,
- Savage, N. (2019). How AI and neuroscience drive each other forwards. *Nature*, 571(7766), 15–17,
- Schuchmann, S. (2019). *History of the Second AI Winter*. (Medium. <https://towardsdatascience.com/history-of-the-second-ai-winter-406f18789d45> (visited on 2023-04-13))
- Sedooka, A., Steffen, G., Paulsen, T., Darbellay, F. (2015). Paradoxe identitaire et interdisciplinarité : un regard sur les identités disciplinaires des chercheurs. *Natures Sciences Sociétés*, 23(4), 367–377,
- Shinn, T., & Joerges, B. (2002). The Transverse Science and Technology Culture: Dynamics and Roles of Research-Technology. *Social Science Information*, 41(2), 207–251,
- Wainer, J., Eckmann, M., Goldenstein, S., Rocha, A. (2013). How productivity and impact differ across computer science subareas. *Communications of the ACM*, 56(8), 67–73,

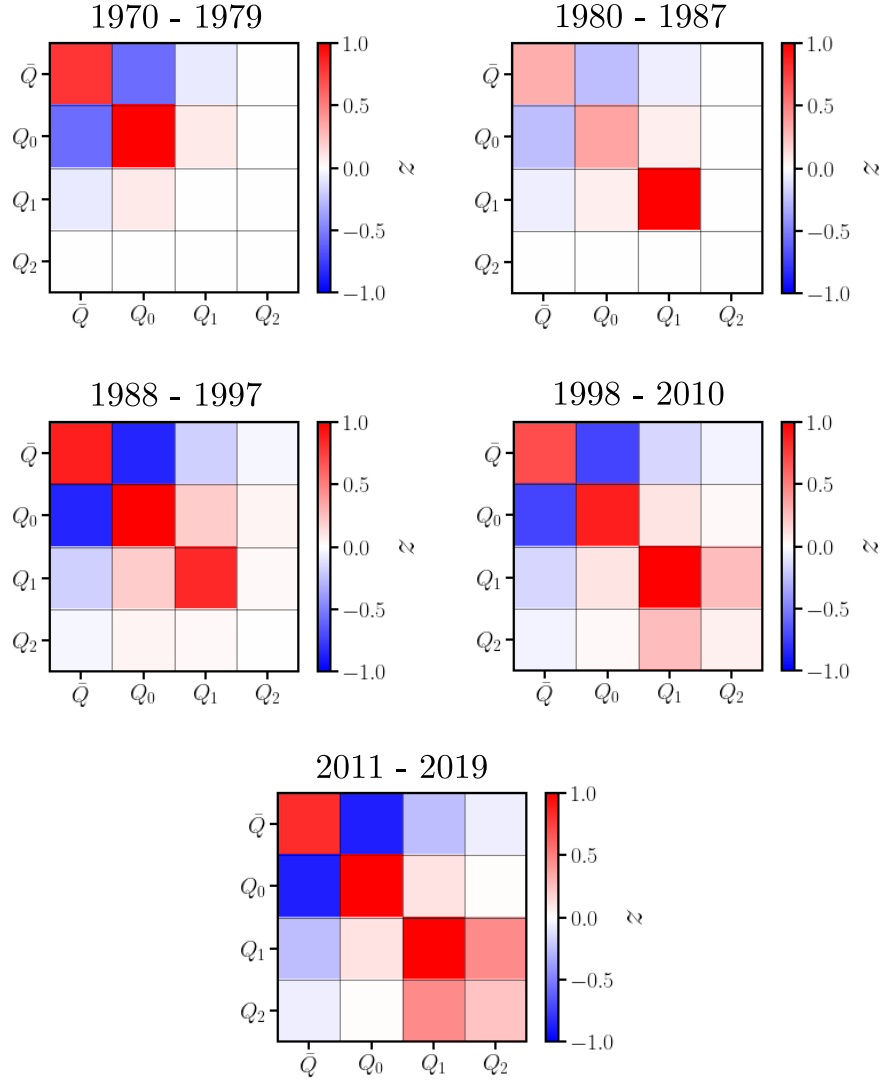


Fig. D3 Evolution of the z-score matrix of the share of edges between f_{AI} quartiles. The first line corresponding to the period 1970-1987 refers to the situation where Q_2 has not yet emerged in our dataset. The separation of \bar{Q} and the other Q_i 's occurs qualitatively in the period 1988-1997 when Q_2 was emerging and was beginning to publish in neuroscience. The ending period spanning from 1998-2010 is very similar to the configuration shown in Fig.13 by considering all the period 1970-2019, as expected by the high concentration of the edges of the temporal collaboration network TCN in this period (91% of the 13,786,616 unweighted single edges in TCN).