v1 – December 16, 2022

The Free Energy Principle & Active Inference: a Systematic Literature Analysis

Virginia Bleu Knight [1,2], RJ Cordes [1,3], Daniel Friedman [1,3,4]

ORCID: VBK: 0000-0002-9894-1989, RJC: 0000-0002-9913-7159, DAF: 0000-0001-6232-9096

Affiliations

1. Active Inference Institute
2. New Mexico State University, Las Cruces NM. Department of Biology
3. Cognitive Security and Education Forum (COGSEC)
4. University of California, Davis. Department of Entomology & Nematology

Keywords

Literature analysis, History of Science, Free Energy Principle, Active Inference, Ontology

Abstract

Here we perform a literature analysis of publications in scientific literature using the term "Free Energy Principle" or "Active Inference", with an emphasis on works written by Karl J Friston. For a subset of papers with accessible full texts, we performed manual annotation (related to structural, visual, and mathematical features) and automated analyses (related to the terms in the Active Inference Institute's Active Inference Ontology). The initial analysis here, at the scale of thousands of citations and hundreds of annotated papers, is presented as a first step towards the development of systems which could:

- Encompass increased scope of relevant works, including non-textual
- Integrate multiple forms of annotation and participation
- Facilitate integration of manual and artificial contributions
- Feature richer interfaces for use in learning & research
- Address field-specific local questions and provide transferable approaches
- Speak to broader questions in the history and philosophy of science

Introduction

One of the primary driving forces of human curiosity is the desire to understand intelligence and consciousness. This motivation has driven the establishment of collective knowledge repositories and the development of an array of algorithms for implementation in cyber-physical systems. Among the mathematical descriptions of learning and decision-making, the Free Energy Principle (FEP) and Active Inference (ActInf) is a recent advancement [1] which builds upon predictive processing and variational inference by incorporating both perception and action into reducing overall uncertainty.

The Free Energy Principle (FEP) and associated process theory of ActInf have evolved markedly over the last 20 years, developing a broad (and at times inchoate) range of discourse. This FEP/ActInf scientific discourse is notable for its range of topics of discussion (e.g. as summarized in the 2022 textbook [1]), and seemingly rapid rate of increase (a pattern which is investigated in this paper). Active Inference is among a handful of other fields that have undergone rapid evolution in the last two decades; other related fields that have had similar transformations include the internet,, cybernetics, and artificial intelligence [2]. In rapidly-changing fields, failures to develop educational and research materials (a situation known as “research debt” [3]) can hamper the growth, accessibility, rigor, and ultimately the value of a scientific development.

Several important textbooks [1,4], reviews [5–8], and retrospectives [9,10] are prominent within the FEP/ActInf field. However these literature reviews and online-only resources [11–13] usually feature narrative or topical reviews or perspectives, highlighting a subset of relevant work (determined coarsely by citation count and public opinion). Conceptual integration across even this restricted set of important works in the field can be challenging, as over the years there are many researchers engaging with different perspectives, amidst a background of theoretical developments. Researchers are engaging with these FEP/ActInf ideas from many different applied, scientific, and philosophical perspectives. Without a proper understanding of the historical and ongoing context of FEP

& ActInf, researchers may waste time engaging with older iterations, hone in on problems that have already been solved, or have increased friction in their attempts to learn and apply these models.

When it comes to technology, the dominant preference is to interact with the most advanced technological products, as well as view evaluations & contextualization of the most recent products. This helps to ensure compatibility with other recent technologies and across current platforms. Similarly, researchers should engage with the latest developments of the FEP, or at the very least, models that are compatible with the current system for optimal learning and applicability. While some overview, narrative, and focused reviews of the FEP/ActInf literature exist [9,14], to our current knowledge there is no empirical meta-analysis aimed at increasing the accessibility and comprehensibility of this increasingly-important body of literature.

We set out to catalog, annotate, and analyze how the FEP and Actinf have developed through time in order to facilitate this emerging area of transdisciplinary and applied engagement. We undertook an analysis of literature related to the Free Energy Principle and ActInf, in an effort to capture and present patterns in publication dynamics (e.g. number of citations, citation networks) and language use (enabled by the Active Inference Ontology, archived at [15]). This work is presented here with minimal narrative and historical analysis, as future analyses will be enabled by improved scope and depth of future iterations of the analysis pipeline.

Methods

Citation discovery & Full text acquisition of open source papers

All code is available in a Knowledge Engineering Github repository, and interactive visualizations are available at the Coda site.

Initially, we used Publish or Perish version 8 [16] for the terms "active inference" and/or "free energy principle" as well as papers authored by Karl J. Friston, from 1990 through 2021, searched with Google Scholar. Results from all the searches were appended into the same dataset, and de-duplicated manually based upon title matching.

To reduce the scope of our initial analysis to open source accessible papers, we used the BioPython API [17] to query for open source publications with titles and/or abstracts containing the terms "active inference" and/or "free energy principle" listed in the NCBI Pubmed Database (Figure 1). Further in-depth analysis and annotation were carried out using this dataset of open source papers.

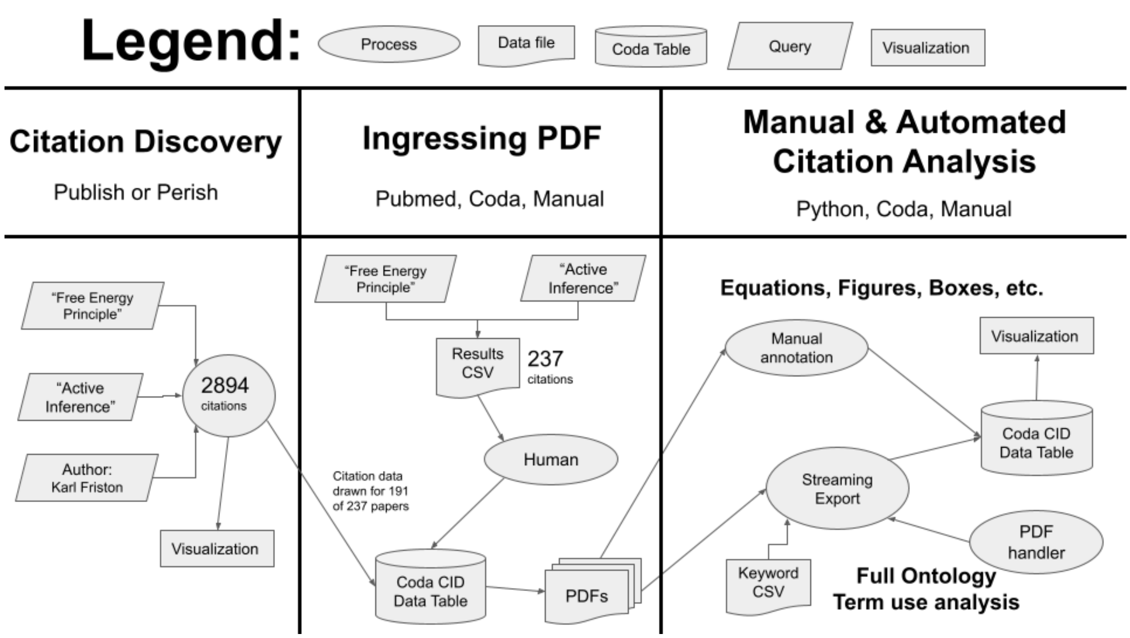

Figure 1. Process overview of the methodology of our pipeline and analysis.

For the set of papers on Pubmed, the available metadata was obtained, including title, abstract, authors, year, and DOI (when available), using the BioPython API and imported using Python and the Pandas library. Citation metrics (from the larger Publish or Perish citation initial analysis) were merged with these data and used to evaluate the number and annual rate of citations for each paper.

PDF files corresponding to each publication were manually downloaded and given a content identifier (CID). Each paper full text PDF corresponded to an entry in an accompanying CSV which contained relevant metadata, including authors, year of publication, and a plain-text abstract. This accompanying CSV was used as an input for the PDF extraction script built specifically for this project (see below).

PDF text extraction

A PDF scraping script was written in Python using the PyPDF2 package [18,19]. PyPDF2 offers simplified text extraction from PDFs of various specifications. Given that the simplified text extraction of PyPDF2 makes no attempt to restructure plain text, the script was written to remove all special characters and spaces in order to form single contiguous blocks before searching for keywords and keyphrases. The script's inputs included a list of ontology terms (74 core terms, 250 supplement terms, and 74 entailed terms, archived at [15]) as a comma separated values file (CSV), a list of PDF files (in CSV format) with a corresponding folder of PDFs, and parameters for where to place results and error logs.

Ontology term frequency analysis

The script first generates a map of document keys which holds all necessary information for accessing a file and temporarily containing its results, and then extracts the ontology terms from each document's plain-text abstract and from the PDF itself using the PyPDF2 library. The results from each PDF are exported before the script moves to the next document, and compiled in a CSV output file containing frequency counts from both the plain-text abstract and PDF file for each document parsed. After initial production of the term frequency data, manual confirmatory analysis using PDF text-search affordances, in comparison with automated keyword counting, did not reveal any strong pattern over- or under-counting of term frequency (see discussion of PDF analysis in the Limitations).

Text annotation

Papers that were analyzed for term usage, were also manually annotated with the following information (Figure 2): # Figures (count of Figures), Figures with equations? (are there equations within Figures), # Equations (written formalisms in the paper), # Tables (number of tables in the paper), # Supplements (number of extra files), # Boxes (number of boxes in the paper), Citation (estimated number of citations), Citations/Year (estimated citations divided by age).

Publications for Analysis

	CID	DOI	pmid	Title	Abstract	First Author	Year	Journal	Authors	# Authors	# of Figures	Figures with Equations?	# Equations	# Tables	# Supplements	# Boxes	Citations	Citations/Year	File
1	00001	10.1007/s00422-014-0620-8	25128318	Active inference, eye movements and oculomotor delays.		Laurent Perrinet	2014	Biological Cybernetics	Laurent Perrinet, Karl J Friston, Rick A Adams	3	10	true	19	0	0	0	46	6.57	
2	00002	10.1007/s00422-018-0753-2	29572721	Planning and navigation as active inference.		Raphael Kaplan	2018	Biological Cybernetics	Raphael Kaplan, Karl J Friston	2	7	true	6	1	0	0	124	41.33	
3	00003	10.1007/s00422-021-00859-9	33471182	Bayesian mechanics of perceptual inference and motor control in the brain.		Chang Sub Kim	2021	Biological Cybernetics	Chang Sub Kim	1	7	false	37	0	0	0	1	1	
4	00004	10.1007/s10339-013-0571-3	23744445	Active inference, sensory attenuation and illusions.		Harriet R Brown	2013	Cognitive Processing	Isabel Parees, Harriet R Brown, Rick A Adams, Karl J Friston, Mark Edwards	5	8	true	4	0	0	0	337	42.13	
5	00005	10.1007/s11023-017-9441-6	31258246	Predictive Processing and the Representation Wars.		Daniel Williams	2018	Minds & Machines	Daniel Williams	1	0	false	0	0	0	0	82	27.33	
6	00006	10.1007/s11229-016-1288-5	29887647	From cognitivism to autopoiesis: towards a computational framework for the embodied mind.		Micah Allen	2018	Synthese	Karl J Friston, Micah Allen	2	4	false	0	0	0	0	287	95.67	
7	00007	10.1007/s11229-021-03075-x	33612868	A critical analysis of Markovian monism.		Majid D Beni	2021	Synthese	Majid D Beni	1	0	false	1	0	0	0	9	9	
8	00008	10.1016/j.bandc.2015.08.002	26275633	The role of interoceptive inference in theory of mind.		Sasha Ondobaka	2017	Brain and Cognition	Sasha Ondobaka, James Kilner, Karl J Friston	3	1	false	0	0	0	0	115	28.75	

Example view of the publications used for analysis. Displayed columns are:

- CID (Content Identifier)
- DOI (Digital Object Identifier)
- pmid (PubMed ID)
- Title (title of the paper)
- Abstract (full text of the Abstract of the paper)
- First Author (name of first author)
- Year (publication year)
- Journal (publication journal)
- Authors (list of all authors)
- # Authors (count of authors)

- # Figures (count of Figures)
- Figures with equations? (are there equations within Figures)
- # Equations (written formalisms in the paper)
- # Tables (number of tables in the paper)
- # Supplements (number of extra files)
- # Boxes (number of boxes in the paper)
- Citation (estimated number of citations)
- Citations/Year (estimated citations divided by age)
- File (link to stored copy of the file)

Figure 2. Manual annotation view of the publications used for analysis. Displayed columns are: CID (Content Identifier), DOI (Digital Object Identifier), pmid (PubMed ID), Title (title of the paper), Abstract (full text of the Abstract of the paper), First Author (name of first author), Year (publication year), Journal (publication journal), Authors (list of all authors), # Authors (count of authors), # Figures (count of Figures), Figures with equations? (are there equations within Figures), # Equations (written formalisms in the paper), # Tables (number of tables in the paper), # Supplements (number of extra files), # Boxes (number of boxes in the paper), Citation (estimated number of citations), Citations/Year (estimated citations divided by age), File (link to stored copy of the file).

Citation network analysis

Research Rabbit [20] was used to analyze citations and construct an co-citation network of the focal set of papers (where nodes represent publications, and edges reflect an instance of one paper citing another). This software produced a JSON file containing all papers in our final dataset as nodes (file available at the code repository), and was visualized in Figure 3.

Coda Implementation

Analysis: The outputs from PubMed, Publish or Perish, PDF analysis, and manual annotation were merged into a single Coda table where each citation was given a content identifier (CID). A list of authors, consisting of all the co-authors of all ingressed papers, was generated. The Active Inference Ontology [21] was downloaded and terms were isolated.

Visualization and Scaling: Coda was used to build reflexive analyses and figures that will be automatically updated as the dataset develops and expands. Additional papers can be added to the dataset, by anyone, by using the form in coda found at the following link [21].

Results

Citation discovery & acquisition

The initial search with Publish or Perish for the terms “active inference” and/or “free energy principle”, or a publication by Karl Friston, resulted in 2894 unique citations (archived file at [15]).

Automated and Manual citation full text annotation

From the large corpus of FEP/ActInf citations, we focused our analysis on an initial set of 237 papers via PubMed. This was the dataset used for subsequent analyses, all of which are presented as an early presentation of a pipeline that can re-render results as new papers and annotations are added (see Limitations).

ResearchRabbit analysis (Figure 3) showed patterns of citation among the papers. No quantitative analysis was performed on this citation network dataset, however the raw output files are provided for future analysis.

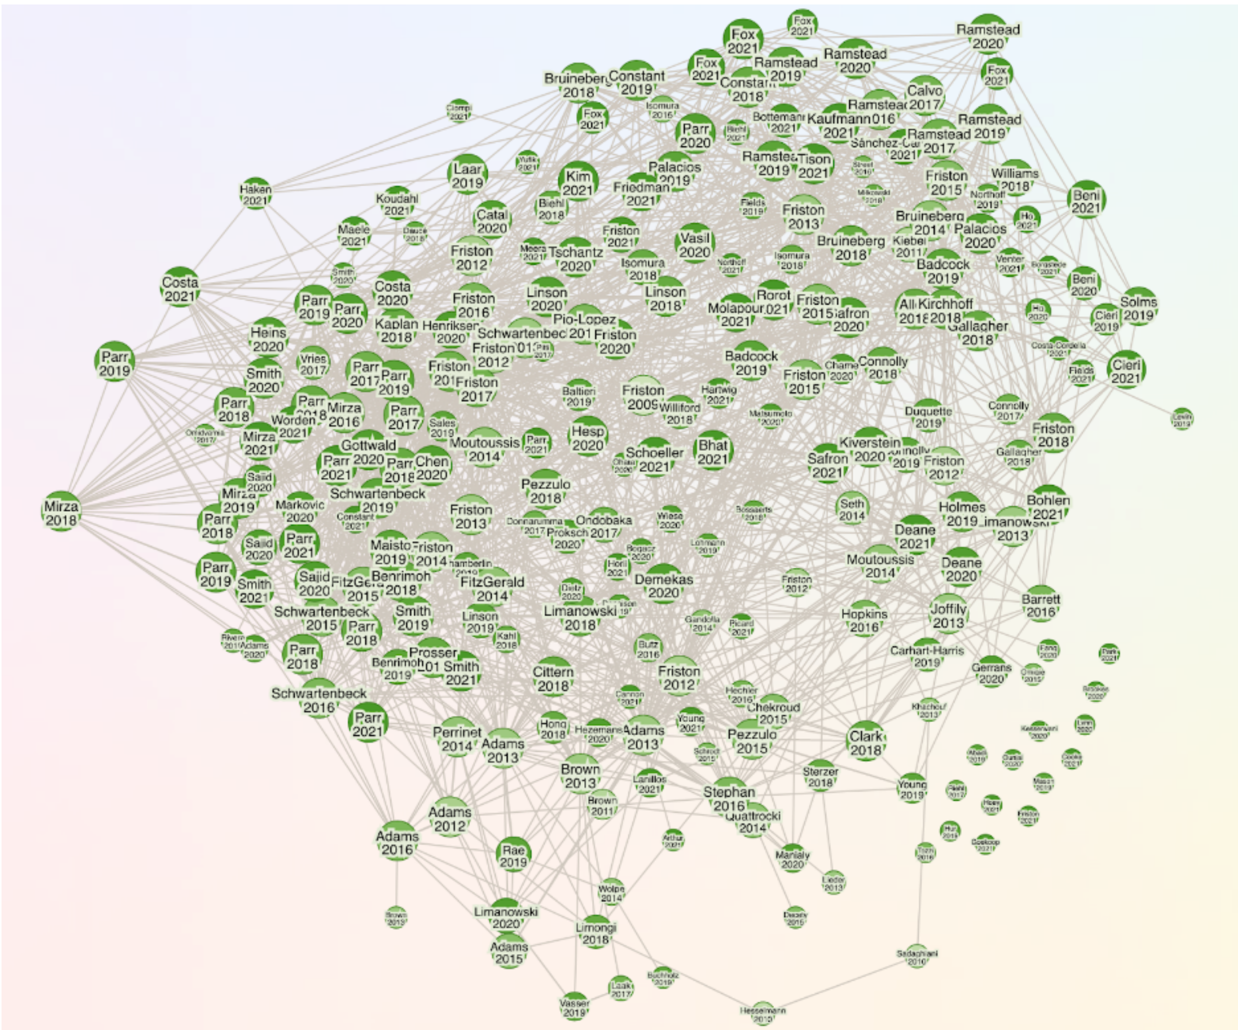

Figure 3. Citation network, by ResearchRabbit. Nodes are papers, identified as (First author, Year). Edges are citations between papers. Data available at Github.

We used our initial, broader citation search in order to identify citation trends among the papers. Figure 4 demonstrates that the most highly cited publications in our open source dataset were from 2013, and all included Karl J. Friston as an author. Shifting our perspective from highly cited papers to highly cited first authors (Figure 5) reveals that the three most highly cited first authors were Karl J. Friston, Rick A. Adams, and Harriet R. Brown.

Citation Heatmap

Last synced 23 hrs ago from here

DOI	Pmid	Year	Authors	Citations	Title	Citations Year
10.1007/s00422-014-0620-8	25128318	2014	Laurent Perrinet, Karl J Friston, Rick A Adams	46	Active inference, eye movements and oculomotor delays.	6.57
10.1007/s00422-018-0753-2	29572721	2018	Raphael Kaplan, Karl J Friston	124	Planning and navigation as active inference.	41.33
10.1007/s00422-021-00859-9	33471182	2021	Chang Sub Kim	1	Bayesian mechanics of perceptual inference and motor control in the brain.	1
10.1007/s10339-013-0571-3	23744445	2013	Isabel Parees, Harriet R Brown, Rick A Adams, Karl J Friston, Mark Edwards	337	Active inference, sensory attenuation and illusions.	42.13
10.1007/s11023-017-9441-6	31258246	2018	Daniel Williams	82	Predictive Processing and the Representation Wars.	27.33
10.1007/s11229-016-1288-5	29887647	2018	Karl J Friston, Micah Allen	287	From cognitivism to autopoiesis: towards a computational framework for the embodied mind.	95.67
10.1007/s11229-021-03075-x	33612868	2021	Majid D Beni	9	A critical analysis of Markovian monism.	9
10.1016/j.bandc.2015.08.002	26275633	2017	Sasha Ondobaka, James Kilner, Karl J Friston	115	The role of interoceptive inference in theory of mind.	28.75

Figure 4: Citation Heatmap. A portion of the coda table found at the Coda site, conditionally formatted based on the number of times each paper has been cited. The number of citations is presented in a rainbow scale, increasing in proportion to the wavelength of the color in the visible spectrum (i.e. purple is the lowest and red is the highest). (a) The first eight rows sorted by citations per year (b) The first six rows sorted by number of citations

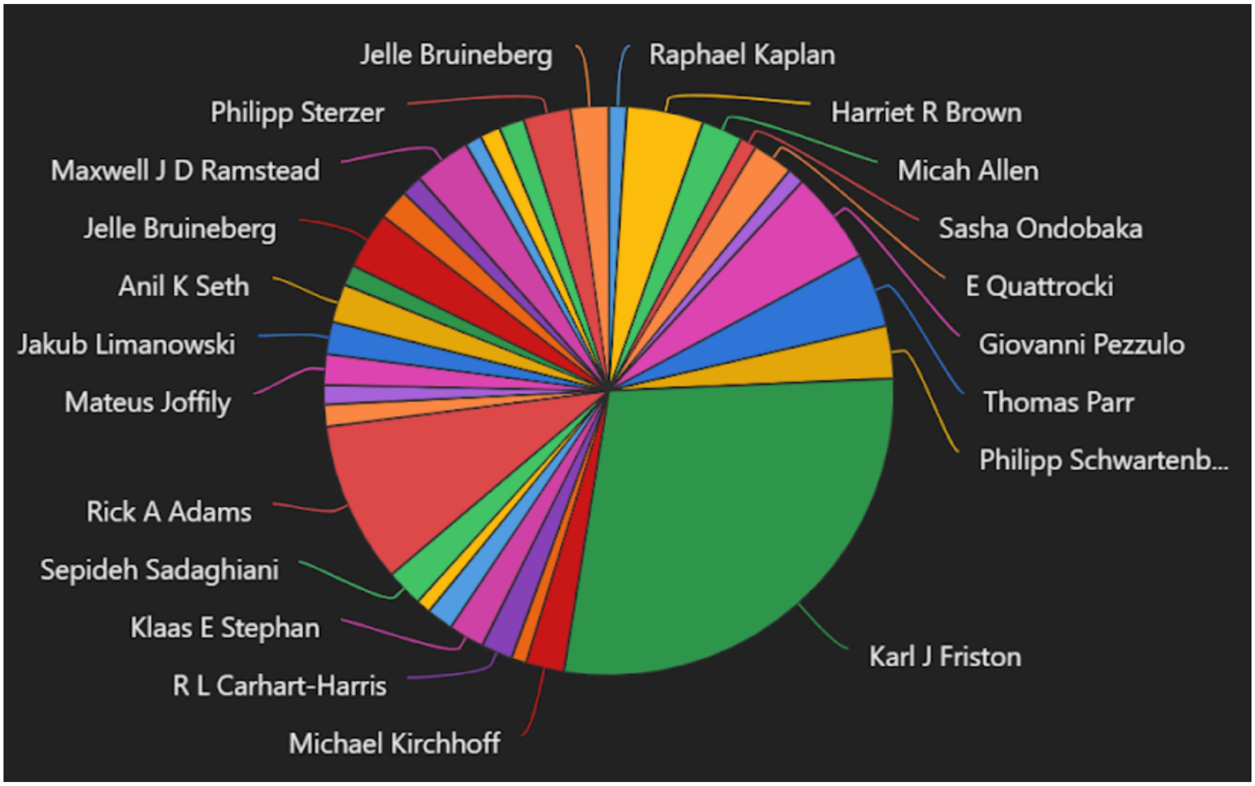

Figure 5: First Author Citations. Citations of open source publications related to active inference and/or the free energy principle, totaled for each first author. The number of citations each first author has received is proportional to the size of the slice of the pie graph. The entire table, including reflexive updates and all underlying data can be found at the Coda site.

Interestingly, when analyzing the number of citations per year for first author papers, Karl J. Friston is not ranked within the top five (Figure 6). The top five first authors with the highest average citations per year are, in order: Phillip Sterzer (109.3), R.L. Carhart-Harris (108.5), Michael Kirchoff (89), Giovanni Pezzulo (71.5), and Maxwell Ramstead (58.2). Of these first authors with the highest number of citations per year, Giovanni Pezzulo was an author of one of the most highly cited ActInf papers of all time (Figure 4b). Three of these first authors (Maxwell Ramstead, Phillip Sterzer, and R.L Carhart-Harris) correspond to publications with the highest number of citations per year (Figure 7). The papers with the highest number of citations per year were all published between 2018-2020 (Figure 7).

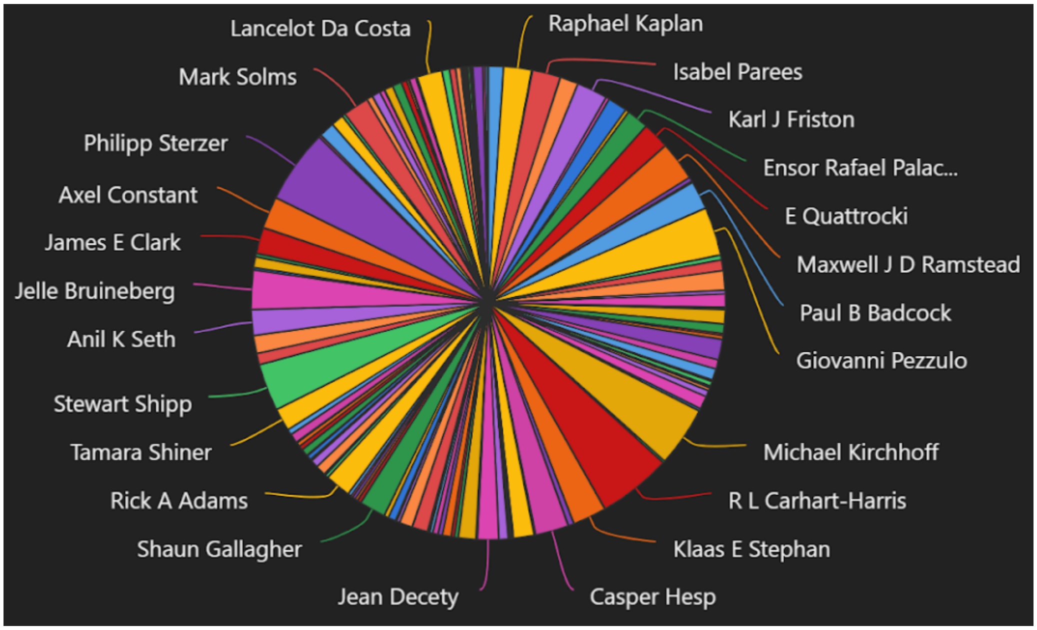

Figure 6: First Author Citations per Year. The average number of citations per year was calculated for each first author. The average number of citations per year for each first author is proportional to the size of the slice of the pie graph. The entire table, including reflexive updates and all underlying data can be found at the Coda site.

A

Year	Citations_per_Year	First_Author	Title
2013	72.63	Karl J Friston	Life as we know it.
2013	70.38	Rick A Adams	The computational anatomy of psychosis.
2013	70.13	Rick A Adams	Predictions not commands: active inference in the motor system.
2012	51.22	Karl J Friston	Perceptions as hypotheses: saccades as experiments.
2009	36.33	Karl J Friston	Reinforcement learning or active inference?
2015	72.33	Giovanni Pezzulo	Active Inference, homeostatic regulation and adaptive behavioural control.

B

Year	Citations_per_Year	First_Author	Title
2020	128	Maxwell J D Ramst	A tale of two densities: active inference is enactive inference.
2018	109.33	Philipp Sterzer	The Predictive Coding Account of Psychosis.
2019	108.5	R L Carhart-Harris	REBUS and the Anarchic Brain: Toward a Unified Model of the Brain Action of Psychedelics.
2019	104	Thomas Parr	Markov blankets, information geometry and stochastic thermodynamics.
2020	97	Casper Hesp	Deeply Felt Affect: The Emergence of Valence in Deep Active Inference.
2018	95.67	Micah Allen	From cognitivism to autopoiesis: towards a computational framework for the embodied mind.

Figure 7: Citations per Year Heatmap. A portion of the Coda site conditionally formatted based on the average number of citations per year. The number of citations per year is presented in a rainbow scale, increasing in proportion to the wavelength of the color in the visible spectrum (i.e. purple is the lowest and red is the highest). (a) The first eight rows sorted by Citations. (b) The first six rows are sorted by Citations per Year

Active Inference Ontology analysis

A frequency map was created which illustrates the sum of the use of all Active Inference ontology terms within each publication (DOI; Figure 8). Analysis of the ontology terms provided a view of term frequency through time (Figure 9). Figure 9A illustrates that all core ontology terms (see Coda site and [22]) increased in use frequency over time. Figure 9B and 10 illustrate how this analysis can be used to drill down to specific terms and identify periods associated with a term's broad adoption in the literature and/ or field of study overall.

Orange Data Mining software [23] was used to visualize and cluster the ontology term use across papers for Figure 10.

Row	Markov Decision Proc	Niche	Selection	Variational	Science	State Space	Stimulus	Error
10.1007/s00422-014-0620-8	0	0	1	7	6	1	20	58
10.1007/s00422-018-0753-2	10	0	9	10	5	9	1	4
10.1007/s00422-021-00859-9	0	5	1	16	4	10	1	24
10.1007/s10339-013-0571-3	0	0	2	1	4	0	13	66
10.1007/s11023-017-9441-6	0	7	0	1	84	0	2	28
10.1007/s11229-016-1288-5	0	11	11	1	45	0	0	19
10.1007/s11229-021-03075-x	0	2	1	4	18	1	1	1
10.1016/j.bandc.2	0	0	0	0	23	0	0	19
10.1016/j.bbr.201	0	0	0	0	0	0	0	0
10.1016/j.heares.	0	0	0	0	0	0	0	0
10.1016/j.jneumet	0	0	2	5	20	1	6	41
10.1016/j.jtbi.201	0	1	11	48	10	2	0	8
10.1016/j.neubior	0	0	10	0	40	0	1	50
10.1016/j.neubior	7	0	18	22	25	0	2	11
10.1016/j.neubior	0	0	0	0	0	0	0	0

Figure 8. Heatmap from the Coda site of Active Inference Ontology term usage. Rows are papers (identified by DOI), columns are terms in the Active Inference Ontology, and cells are red proportionately to use of that term in the paper.

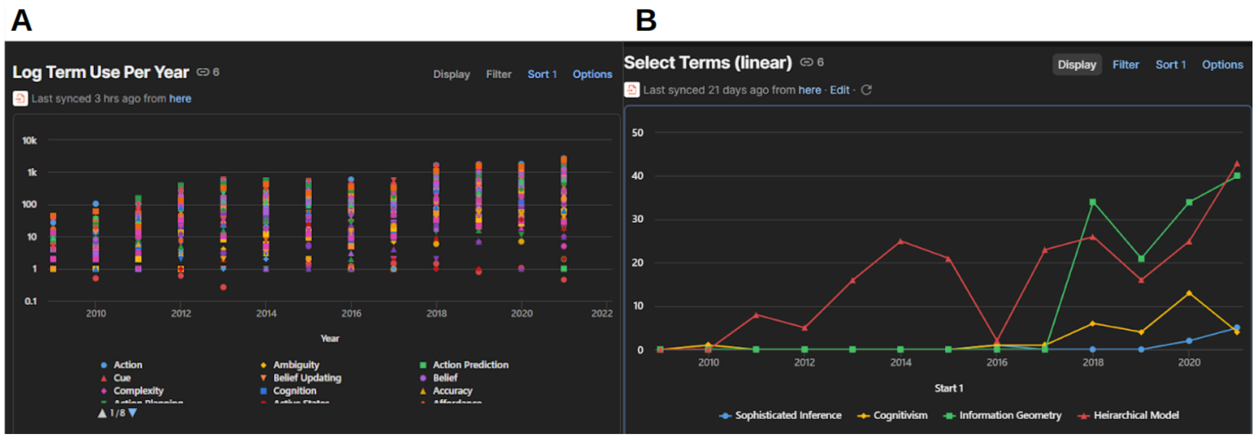

Figure 9. Term use over time. The terms from the active inference ontology were cataloged within each document and summarized across all documents in the open source dataset. (a) All terms from the Active Inference Institute core ontology plotted by frequency and year on a logarithmic scale (b) Select core ontology terms plotted by frequency and year on a linear scale. These terms were selected to evaluate in more detail because of their relevance to highly cited publications (Figure J). The entire table, including reflexive updates and all underlying data can be found at the Coda site.

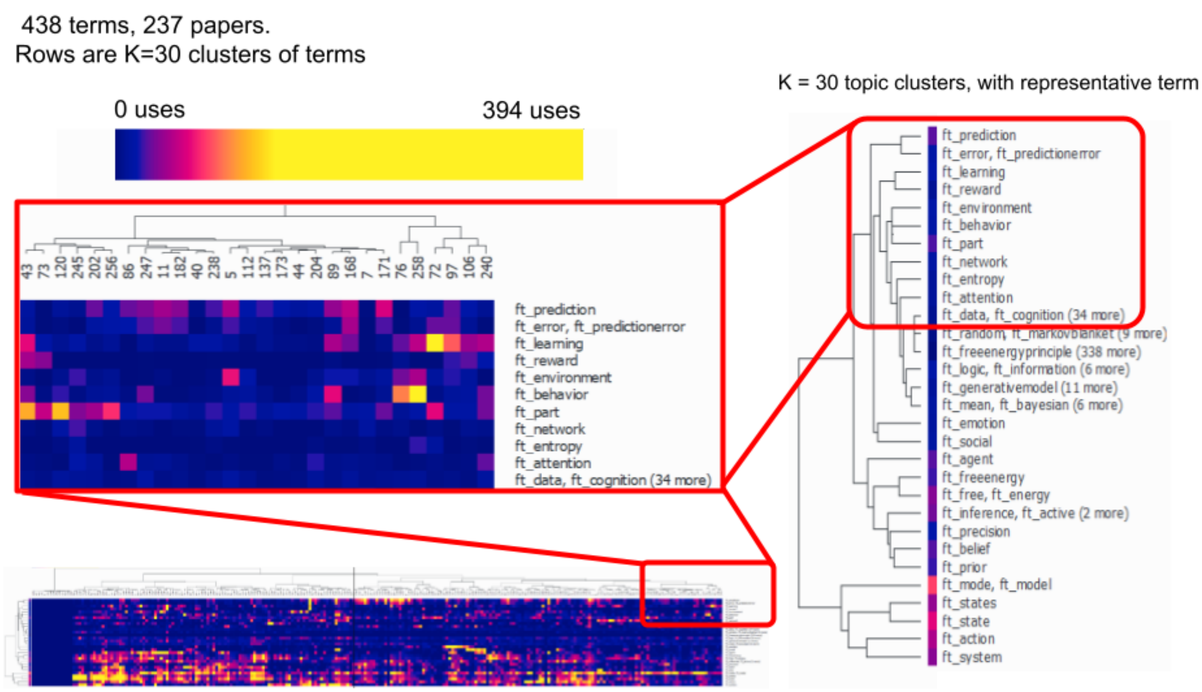

Figure 10. Heatmap of Ontology terms. Columns are papers (identified by content ID), rows are single or clustered ontology terms (K=30 clustering). Cells are colored according to their use in the paper, ranging from zero uses (purple) to 394 uses (yellow).

Discussion

Summary

Authorship Trends

Karl J. Friston is the center figure in the set of papers analyzed here, in terms of number of papers and citation intensity. Overall, many of the most highly cited first authors (Figure 5) correspond to the most highly cited papers (Figure 4) in our dataset. When evaluating authorship by the number of first author citations per year, the landscape shifts dramatically (Figure 6). While Professor Friston's first author papers have enough citations per year to put him on the chart, this number is surpassed by the citations per year for P. Sterzer, G. Pezzulo, M. Kirchoff, R.L. Carhart-Harris, and others.

Citation Trends

As expected, the top six most cited papers in our open-source dataset all included Karl J. Friston as an author (Figure 4). These papers included, in order: "Life as we know it" (Friston 2013) [24], "The computational anatomy of psychosis" (Adams et al., 2013a) [25], "Predictions not commands: active inference in the motor system" (Adams et al., 2013b) [26], "Perceptions as hypotheses: saccades as experiments" (Friston et al., 2012) [27], "Reinforcement learning or active inference?" (Friston et al., 2009) [28], and "Active inference, homeostatic regulation, and adaptive behavioural control" (Pezzulo et al., 2015) [29]. Half of these papers were from 2013, which corresponds to the beginning of something analogous to a Cambrian explosion in the field of ActInf. Since that time, we have seen the number of papers increase exponentially (Figure 11). The frequency of the Active Inference Institute core ontology terms has also increased exponentially through the years since 2013 (Figure 9). We can infer that the increase in the frequency of ontology term use over time corresponds to the increased numbers of relevant papers through time (Figure 11).

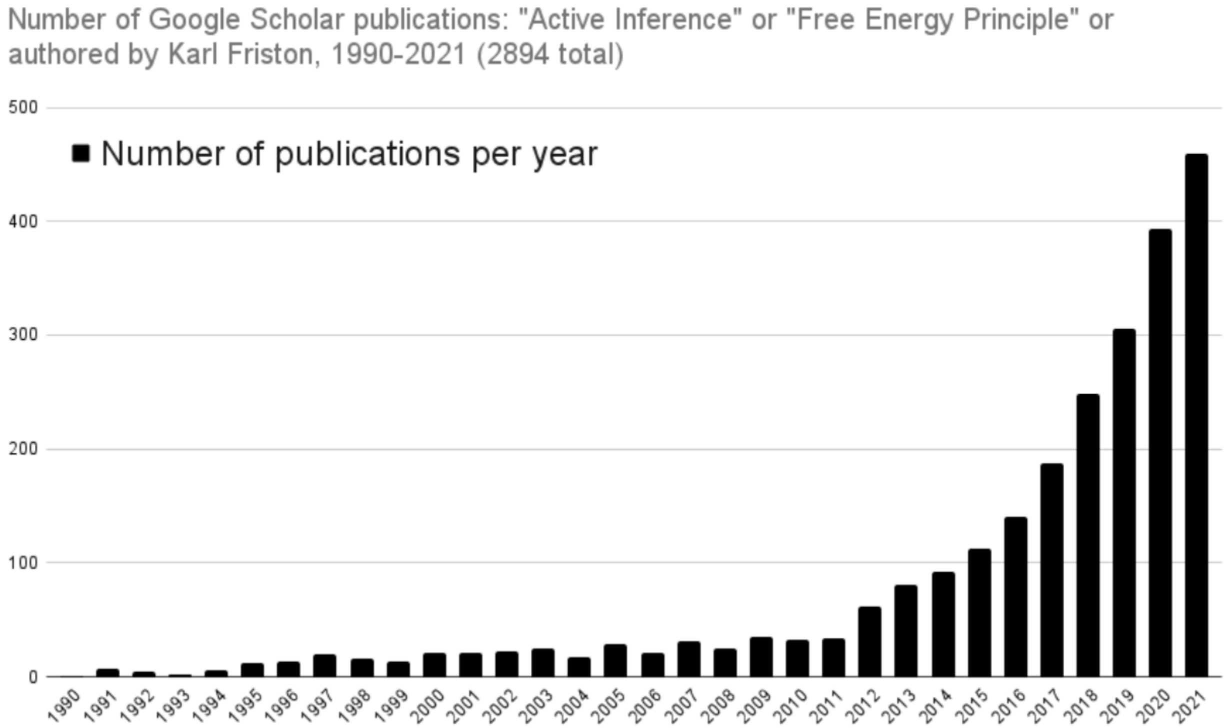

Figure 11. Number of papers published per year (indexed on Google Scholar) mentioning “Active Inference” or “Free Energy Principle”, or with Karl Friston as co-author.

The most highly cited papers in our dataset are all seminal works in the field. The work of Friston et al. (2009) develops an adaptive agent that can solve a dynamic programming benchmark- the mountain-car problem, using ActInf under the FEP (Friston et al., 2009). Friston et al. (2012) model eye saccades as experiments that gather data which test hypotheses about the underlying causes of the data, thus demonstrating that visual search can be motivated by minimizing the entropy of the world’s hidden states. In the most highly cited work in our dataset, Friston (2013) provides evidence that homeostasis and autopoiesis are properties of ergodic random dynamical systems possessing a Markov blanket, which renders internal states conditionally independent from external states (Friston 2013). The work of Adams et al. (2013a) models psychotic symptoms as a deviation from normal, Bayes-optimal action and perception- specifically as a reduction in the precision of beliefs relative to sensory evidence. In the work of Adams et al. (2013b), motor behavior is modeled under the framework of ActInf, with the motor cortex sending proprioceptive predictions as opposed to motor commands. The work of Pezzulo et al. (2015) unifies the behavioral theory of ActInf- a Bayesian, probabilistic formulation of perception and action- with the many associative learning (Pavlovian, habitual, goal-directed) theories that underlie homeostatic and behavioral control (Pezzulo et al., 2015)

Examining the citation trends through the lens of citation number per year provides an alternative view of heavily cited papers by placing newer articles on equal footing with many of the seminal articles in the field. The six papers in our dataset with the highest number of citations per year (Figure 7) include, in order: “A tale of two densities: active inference is enactive inference” (Ramstead et al., 2020) [30], “The predictive coding account of psychosis” (Sterzer et al., 2018) [31], “Rebus and the anarchic brain: towards a unified model of the brain action of psychedelics” (Carhart-Harris & Friston, 2019) [32], “Markov blankets, information geometry, and stochastic thermodynamics” (Parr et al., 2019) [33], “Deeply felt affect: the emergence of valence in deep active inference” (Hesp et al., 2020) [34], “From cognitivism to autopoiesis: a computational framework for the embodied mind” (Friston & Allen, 2018) [35]. In these highly cited papers, Karl J. Friston is an author on all except one: “The predictive coding account of psychosis.” Whereas the set of papers with the most citations largely comprised early work in the field, the papers with the most citations per year in our dataset were all written between 2018 and 2020.

These recent highly cited works are critical in the field for different reasons. Ramstead et al. (2020) clarified how generative models and variational densities had been largely misinterpreted in the literature because of the confusion between ActInf and the FEP, and alternative Bayesian frameworks such as predictive coding. Sterzer et al. (2018) relate the predictive coding mechanism of brain function to the aberrant neurotransmission and phenomenology that is seen in psychosis. Carhart-Harris & Friston (2019) integrate the entropic brain hypothesis and the FEP to account for the consciousness-altering action of psychedelics with a unified model of altered brain activity. Parr et al. (2019) posit the information geometry of the Markov blanket as the foundation for describing the relationships between inference, information, and thermodynamics. Hesp et al. (2020) provides a unified account of mental action, affect, and metacognition, which is founded on the scaffold of deep active inference. Friston & Allen (2018) review predictive processing and identify a schism between enactive and embodied approaches, which they begin to unify through the FEP.

Ontology Trends

The salience of the highly cited papers identified above demonstrates that meta-analyses can identify key pivot-points in nascent fields such as ActInf. Similarly, the evaluation of term frequency can also be leveraged to identify trends and critical changes in the field through time. For example, Figure 9B highlights the frequency of some selected terms corresponding to the highly cited publications from 2018-2020, described above. Here it can be seen that the frequency of the term Information Geometry increased following the publication by Parr et al., 2019 which focused on the significance of the information geometry of the Markov blanket. Similarly, use of the term Cognitivism in the active inference literature has increased since the related publication by Friston & Allen (2018). Furthermore, the Hesp et al., 2020 publication leveraged the Hierarchical Model of Sophisticated Inference, and both of these terms increased in frequency leading up to and following the release of this publication.

Limitations & Future approaches

Citation discovery & acquisition

Data availability was limited for Publish or Perish (citation) as well as PubMed (citation). For example, several open source papers from the open source PubMed dataset were not listed in the Publish or Perish data, despite identical terms used for the search. PubMed data did not always contain the DOI. Therefore, there was a great degree of manual curation for these data. Many papers in the FEP and ActInf literature are not available open source. The open source restriction in the PubMed query reduced the number of papers from ~3000 in the Publish or Perish dataset to ~220 papers that were pushed to the analysis pipeline.

The pipeline and online data will continue to be stewarded by the Active Inference Institute, and we expect to update the ontology and online figures continually, with volunteer and employee efforts. We will simplify PDF upload and metadata annotation through the online software Coda, and it is our hope that authors who wish to have their data included in future renditions of these analyses will upload their articles to our database.

Manual annotation

Current manual annotation of paper structure (e.g. number of Tables, Figures, Equations, Supplemental files) was presented as preliminary, in recognition of the many challenges facing the development of reliable and accessible crowdsourcing systems [36–38]. Future iterations of this system could incorporate best practices in distributed scholarship, to facilitate the addition, verification, integration, and presentation of massive annotation databases. Already however, the initial manual annotation provides helpful functionality to the seeker by for example, allowing one to restrict their focus only to papers which do or do not have equations (including equations within figures, a common style in the ActInf/FEP literature considered).

PDF text

The Portable Document Format (PDF) is styled and rendered for display, and is not always amenable for conversion to plain-text. This creates numerous information quality complications in bibliometric analysis. For example, some papers could not be scraped with our existing methods and therefore the number of ontology terms isn't listed in the outcomes. Not all documents could be guaranteed to include the contents of their plain text abstract verbatim in their corresponding PDF, nor can it be guaranteed that, if they did, keyword extraction would be free of false negatives (due to e.g. a hyphenated word), thus, abstract count and document count were kept as separate metrics. Due to the difficulties communicated above, we cannot be completely certain of accurate keyword extraction.

Despite the PDF's central role as standardized file type used for academic publishing, PDFs are notably difficult to computationally parse. This challenge has implications for the accessibility, rigor, and security of research contained in PDFs [39–41]. The PDF specification has myriad versions and complexities, and, given that these specifications can contain hundreds of pages of details, not all PDF implementations are strictly following the specifications of the version they are marked as [42]. Even where specifications are strictly followed, there are still serious challenges due to the presence of multiple forms of difficult to extract or difficult to reconstruct text elements, including, but not limited to, page numbers, page headers and footers, text within tables, hidden text, form inputs, text within images, paragraph text placed as images, image captions, noncontiguous paragraph text, and ligatures [43]. Further, the fact that PDFs are (generally) used to render styled and subjective content and that there are numerous use-cases for extraction means that there cannot be a general computational method for verification of what constitutes "correct" output to linear plain-text. Standardization of publication styling and format could theoretically address these problems for future publications, though that standardization may adversely impact the aesthetics (and content) of creators and publishers and adoption would represent its own challenges. We suggest that affordances for annotation and related standards may provide opportunities to improve the time-to-impact and reduce complications of bibliometric study [44,45].

Ontology term analysis

The current ontology contains nested terms, for example "Free Energy Principle" and "Free Energy". Based on the simple syntactic system we implemented for this work, the count for "Free Energy" also contains "Free Energy Principle". In many articles, terms are abbreviated with an acronym, such as FEP for Free Energy Principle. This analysis does not account for acronym usage. In the ActInf and FEP literature, multiple terms are often used by authors to denote the same thing (i.e. observation, outcome, sensory outcome). The current ontology typically only accounts for one of the possible terms. Future renditions of the ontology (and subsequent literature analysis) should account for possible meaning overlap between terms, as well as acronym usage, alternative spellings (e.g. American and British English), and non-English languages.

Next steps

Some specific next steps are described above, in relation to the Limitations they address. Some further directions of development include:

- Analysis of per-author and overall linguistic patterns, to identify trends and changepoints in language use (term frequency, co-occurrence, etc).
- The citation network analysis presented could be expanded in scope (e.g. including more citations), and connected more closely with linguistic analyses. For example, what are the patterns of language use in papers that are citation network outliers (e.g. with more or less connectivity to other ActInf citations)?
- Development of practices around using the results of bibliographic analyses, in daily education and research settings.

Conclusion

The evolution of fields such as mathematics, biology, and others demonstrates how the significant findings of prominent individuals have had a substantial historical impact. Looking back, we can see how these contributions introduced new methods and vocabulary across deep time. Computational meta-analyses, such as the work presented here, allow us to track author impacts and vocabulary changes in real time, increasing the ability to sensemake and act in a rapidly developing epistemic niche. The ongoing stewardship of this project should offer continuity and facilitate the deployment of rigorous, accessible research products and interfaces going forward. Voluntary contributions to data for this pipeline are encouraged via submission through the form at the following Coda link.

Code Availability & Maintenance

All code is available in the Active Inference Institute Github repository titled "Knowledge Engineering" (https://github.com/ActiveInferenceInstitute/Knowledge-Engineering [15]). The stewards of this pipeline and project from 2023 on, will be the Active Inference Institute.

References

1. Parr T, Pezzulo G, Friston KJ. Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. 2022.

2. Cordes RJ, Friedman D. “The Great Preset: Remote Teams and Operational Art.” 2020 [cited 19 Oct 2021]. Available: https://www.cogsec.org/research-irt

3. Olah C, Carter S. Research Debt. Distill. 2017;2. doi:10.23915/distill.00005

4. Friston K. Statistical parametric mapping. Statistical Parametric Mapping. 2007. pp. 10–31. doi:10.1016/b978-012372560-8/50002-4

5. Friston K. The free-energy principle: a unified brain theory? Nat Rev Neurosci. 2010;11: 127–138. doi:10.1038/nrn2787

6. Sajid N, Ball PJ, Parr T, Friston KJ. Active inference: demystified and compared. arXiv [cs.AI]. 2019. Available: http://arxiv.org/abs/1909.10863

7. Da Costa L, Parr T, Sajid N, Veselic S, Neacsu V, Friston K. Active inference on discrete state-spaces: A synthesis. J Math Psychol. 2020;99: 102447. doi:10.1016/j.jmp.2020.102447

8. Millidge B, Seth A, Buckley CL. Predictive Coding: a Theoretical and Experimental Review. arXiv [cs.AI]. 2021. Available: http://arxiv.org/abs/2107.12979

9. Friston K. The history of the future of the Bayesian brain. Neuroimage. 2012;62: 1230–1233. doi:10.1016/j.neuroimage.2011.10.004

10. Friston K, Fortier M, Friedman DA. Of woodlice and men. ALIUS Bulletin. 2018;2: 17. Available: https://www.aliusresearch.org/uploads/9/1/6/0/91600416/alius_bulletin_n%C2%B02__2018_.pdf#page=27

11. Beren. FEP_Active_Inference_Papers: A repository for major/influential FEP and active inference papers. Github; 2021. Available: https://github.com/BerenMillidge/FEP_Active_Inference_Papers

12. Tumiel J. Spinning up in active inference and the free energy principle. In: jared tumiel [Internet]. 14 Oct 2020 [cited 21 Jan 2022]. Available: https://jaredtumiel.github.io/blog/2020/10/14/spinning-up-in-ai.html

13. Active Inference Institute. In: Active Inference Institute, homepage [Internet]. [cited 15 Feb 2022]. Available: https://www.activeinference.org/

14. Andrews M. The math is not the territory: navigating the free energy principle. Biol Philos. 2021;36: 30. doi:10.1007/s10539-021-09807-0

15. Active Inference Institute. Knowledge-Engineering Github repository. Github; Available: https://github.com/ActiveInferenceInstitute/Knowledge-Engineering

16. Faiman B. Publish or Perish. J Adv Pract Oncol. 2021;12: 463–464. doi:10.6004/jadpro.2021.12.5.1

17. Cock PJA, Antao T, Chang JT, Chapman BA, Cox CJ, Dalke A, et al. Biopython: freely available Python tools for computational molecular biology and bioinformatics. Bioinformatics. 2009;25: 1422–1423. doi:10.1093/bioinformatics/btp163

18. PyPDF2. Github; 2022. Available: https://github.com/py-pdf/PyPDF2

19. Cordes RJ. Half22_PDF_Scrape. Github; 2022. Available: https://github.com/ActiveInferenceLab/Knowledge-Engineering

20. Human Intelligence Technologies, Incorporated. ResearchRabbit (Version 4.0). ResearchRabbit. Available: https://www.researchrabbit.ai/

21. Active Inference Institute. Active inference ontology. In: Coda | Everything evolves, the evolution of documents [Internet]. 7 Mar 2022 [cited 12 Dec 2022]. Available: https://coda.io/@active-inference-institute/active-inference-ontology-website

22. Active Inference Institute. ActiveInferenceInstitute/Active_Inference_Ontology: December 12, 2022 Snapshot. 2022. doi:10.5281/zenodo.7430333

23. Bioinformatics Laboratory, University of Ljubljana. Orange Data Mining. [cited 16 Dec 2022]. Available: https://orangedatamining.com/

24. Friston K. Life as we know it. J R Soc Interface. 2013;10: 20130475. doi:10.1098/rsif.2013.0475

25. Adams RA, Stephan KE, Brown HR, Frith CD, Friston KJ. The computational anatomy of psychosis. Front Psychiatry. 2013;4: 47. doi:10.3389/fpsyt.2013.00047

26. Adams RA, Shipp S, Friston KJ. Predictions not commands: active inference in the motor system. Brain Struct Funct. 2013;218: 611–643. doi:10.1007/s00429-012-0475-5

27. Friston K, Adams RA, Perrinet L, Breakspear M. Perceptions as hypotheses: saccades as experiments. Front Psychol. 2012;3: 151. doi:10.3389/fpsyg.2012.00151

28. Friston KJ, Daunizeau J, Kiebel SJ. Reinforcement Learning or Active Inference? PLoS ONE. 2009. p. e6421. doi:10.1371/journal.pone.0006421

29. Pezzulo G, Rigoli F, Friston K. Active Inference, homeostatic regulation and adaptive behavioural control. Prog Neurobiol. 2015;134: 17–35. doi:10.1016/j.pneurobio.2015.09.001

30. Ramstead MJD, Kirchhoff MD, Friston KJ. A tale of two densities: active inference is enactive inference. Adapt Behav. 2019;28: 1059712319862774. doi:10.1177/1059712319862774

31. Sterzer P, Adams RA, Fletcher P, Frith C, Lawrie SM, Muckli L, et al. The Predictive Coding Account of Psychosis. Biol Psychiatry. 2018;84: 634–643. doi:10.1016/j.biopsych.2018.05.015

32. Carhart-Harris RL, Friston KJ. REBUS and the Anarchic Brain: Toward a Unified Model of the Brain Action of Psychedelics. Pharmacol Rev. 2019;71: 316–344. doi:10.1124/pr.118.017160

33. Parr T, Da Costa L, Friston K. Markov blankets, information geometry and stochastic thermodynamics. Philos Trans A Math Phys Eng Sci. 2020;378: 20190159. doi:10.1098/rsta.2019.0159

34. Hesp C, Smith R, Parr T, Allen M, Friston K, Ramstead M. Deeply Felt Affect: The Emergence of Valence in Deep Active Inference. 2019. doi:10.31234/osf.io/62pfd

35. Allen M, Friston KJ. From cognitivism to autopoiesis: towards a computational framework for the embodied mind. Synthese. 2018;195: 2459–2482. doi:10.1007/s11229-016-1288-5

36. Sabou M, Bontcheva K, Derczynski L, Scharl A. Corpus Annotation through Crowdsourcing: Towards Best Practice Guidelines. Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14). Reykjavik, Iceland: European Language Resources Association (ELRA); 2014. pp. 859–866. Available: http://www.lrec-conf.org/proceedings/lrec2014/pdf/497_Paper.pdf

37. Xia H, Østerlund C, McKernan B, Folkestad J, Rossini P, Boichak O, et al. TRACE: A Stigmergic Crowdsourcing Platform for Intelligence Analysis. Proceedings of the 52nd

Hawaii International Conference on System Sciences. scholarspace.manoa.hawaii.edu; 2019. Available: https://scholarspace.manoa.hawaii.edu/handle/10125/59484

38. Kim S, Robert LP Jr. Crowdsourcing Coordination: A Review and Research Agenda for Crowdsourcing. Macrotask Crowdsourcing: Engaging the Crowds to Address Complex Problems. 2019. Available: https://books.google.com/books?hl=en&lr=&id=f8anDwAAQBAJ&oi=fnd&pg=PA17&dq=stigmergy+complexity+ant&ots=Sz8nwnp4Vx&sig=_AOF0KGJcpljuBZT1n3E_Z5aVpQ

39. Adobe Systems. PDF Reference: Adobe Portable Document Format Version 1.4. Addison-Wesley; 2001. Available: https://play.google.com/store/books/details?id=ijFdTZDQugwC

40. Castiglione A, De Santis A, Soriente C. Security and privacy issues in the Portable Document Format. J Syst Softw. 2010;83: 1813–1822. doi:10.1016/j.jss.2010.04.062

41. Rautiainen S. A look at Portable Document Format vulnerabilities. Information Security Technical Report. 2009;14: 30–33. doi:10.1016/j.istr.2009.04.001

42. Thoma M. PYPDF2 - Robustness And Strict = False. Available: https://github.com/py-pdf/PyPDF2/blob/39215c704791c10189668bbd1fa1d04d0b1f3f81/docs/user/robustness.md

43. Thoma M. PYPDF2 - Extract Text from a PDF. Available: https://github.com/py-pdf/PyPDF2/blob/39215c704791c10189668bbd1fa1d04d0b1f3f81/docs/user/extract-text.md

44. David S, Cordes RJ, Friedman DA, editors. Structuring the Information Commons: Open Standards and Cognitive Security. COGSEC; 2022.

45. Tamari R, Friedman D, Fischer W, Hebert L, Shahaf D. From Users to (Sense)Makers: On the Pivotal Role of Stigmergic Social Annotation in the Quest for Collective Sensemaking. Proceedings of the 33rd ACM Conference on Hypertext and Social Media. New York, NY, USA: Association for Computing Machinery; 2022. pp. 236–239. doi:10.1145/3511095.3536361