Biology & Philosophy (2020) 35:5
https://doi.org/10.1007/s10539-019-9729-4

Revisiting three decades of *Biology and Philosophy*: a computational topic-modeling perspective

Christophe Malaterre[1,2] · Davide Pulizzotto[3] · Francis Lareau[4]

Received: 20 May 2019 / Accepted: 29 November 2019 / Published online: 19 December 2019

Abstract
Though only established as a discipline since the 1970s, philosophy of biology has already triggered investigations about its own history (e.g., Grene and Depew in The philosophy of biology: an episodic history, Cambridge University Press, Cambridge, 2004. https://doi.org/10.1017/CBO9780511819018; Hull, in: Ruse (ed) The Oxford handbook of philosophy of biology, Oxford University Press, New York, pp 11–33, 2008). When it comes to assessing the road since travelled—the research questions that have been pursued—manuals and ontologies also offer specific viewpoints, highlighting dedicated domains of inquiry and select work. In this article, we propose to approach the history of the philosophy of biology with a complementary data-driven perspective that makes use of statistical algorithms applied to the complete full-text corpus of one major journal of the field—*Biology and Philosophy*—from its launch in 1986 up until 2017. By running text-mining and topic-modeling algorithms, we identified 67 key research topics that span across these 32 years. We also investigated the evolution of these topics over time and their fluctuating significance in the journal articles. Our results concur with known episodes or traits of the discipline—for instance, the significance of evolution-related topics or the decrease of articles with a marked historical dimension—but also highlight a diversity of topics that is much richer than what is usually acknowledged.

Keywords Biology and Philosophy journal · History of philosophy of science · Topic modeling · Text mining · Digital history · Digital philosophy

Electronic supplementary material The online version of this article (https://doi.org/10.1007/s10539-019-9729-4) contains supplementary material, which is available to authorized users.

✉ Christophe Malaterre
malaterre.christophe@uqam.ca

Extended author information available on the last page of the article

Introduction

In the past few decades, the philosophy of biology has grown into a specific subdiscipline of philosophy and philosophy of science, with its own learned society and dedicated journals. This is not to say that no philosophy of biology was done before: only that such work was not considered as constituting a recognizable subset within established philosophical traditions, with its dedicated set of authors and scholarly institutions. Large-scale histories of the philosophy of biology often go back to Aristotle, survey selected works of Descartes and Harvey, Buffon and Kant, describe the influential role of eighteenth–nineteenth century philosophers and scientists—among which Herschel, Whewell, Lyell, Mill, Comte but also Jevons, Pierce, James or Dewey—and position the rise of a specific philosophy of biology in the second half of the twentieth century, with the early works of Woodger, Grene and Beckner in the late 1950s, the momentum created by Hull, Ruse, Schaffner and Wimsatt in the 1970s, culminating in the institutionalization of the discipline in the 1980s (Ruse 1989; Callebaut 1993; Grene and Depew 2004; Byron 2007; Hull 2008; Gayon 2009; Nicholson and Gawne 2015; Ruse 2016).[1] One of the leading journals of the field—*Biology and Philosophy*—indeed launched in 1986, and the International Society for the History, Philosophy and Social Studies of Biology (ISHPSSB) was formed a few years later in 1989, both constituting tangible signs of the professionalization of the discipline.

The recent history of the philosophy of biology and, even more interestingly, of its research questions—especially as they compare to those tackled by the philosophy of science in general—is reflected in the contents of an increasing number of dedicated textbooks (Ruse 1973; Hull 1974; Ruse 1988; Sober 1999; Sterelny and Griffiths 1999; Garvey 2007; Rosenberg and McShea 2008; Godfrey-Smith 2014), edited volumes that offer review chapters on specific topics (Ayala and Dobzhansky 1974; Hull and Ruse 2007; Matthen and Stephens 2007; Ruse 2008; Sarkar and Plutynski 2008; Ayala and Arp 2010; Kampourakis 2013), and anthologies that showcase selected foundational articles (Grene and Mendelsohn 1976; Hull and Ruse 1998; Sober 2006; Rosenberg and Arp 2009). In this article, we offer a complementary perspective on the philosophy of biology that makes use of computational text-mining methodologies. By applying topic-modeling algorithms to the full-text corpus of *Biology and Philosophy*, we investigated the major research themes that mobilized philosophers of biology in the last thirty or so years, and we quantitatively analyzed how these themes evolved over this period of time. Whereas previous analyses of the journal publications have been carried out by manually sorting publications into pre-defined categories (Gayon 2009; Pradeu 2017), we relied on

[1] Though philosophy of biology was not established as a specific discipline in the first half of the twentieth century, a reasonable share of articles published in general philosophy of science journals at that time did concern biology-related topics (see (Byron 2007; Nicholson and Gawne 2015); see also (Malaterre et al. 2019) for a text-mining perspective). These early works are often neglected in commonly accepted narratives about the history of the discipline (e.g., Callebaut 1993; Hull 2008), likely due to a perspective stressing the rarity of individuals doing good work in the field, their lack of knowledge of the biological literature and the absence of institutional support at that time (Hull 1969; Honenberger 2018, 296 n. 28).

computational approaches that could scan the entire corpus to generate topics in a bottom-up fashion, based on the semantic content of the publications themselves. These approaches indeed make it possible to analyze more comprehensively and more systematically the contents of large corpora of full-text documents, be it in history and sociology (e.g., Chartier and Meunier 2011; Mimno 2012; DiMaggio et al. 2013; Evans and Aceves 2016; Peirson et al. 2017; Barron et al. 2018), linguistics and the cognitive sciences (e.g., Turney and Pantel 2010; Widdows 2004; Murdock et al. 2017), or philosophy (Buckner et al. 2011; Ramsey and Pence 2016; Pence and Ramsey 2018; Malaterre et al. 2019). The article is therefore not a traditional philosophical paper, yet we believe it should be of interest to many readers of *Biology and Philosophy*—experts and novices alike—and we hope enjoyable. The data we are presenting are the type of data that offer an empirical basis for what might otherwise be informal claims about the discipline and its evolution in the past decades. They are also the type of data that may prompt discussion about the directions the field might—or could—take.

The paper is organized as follows. We first describe the corpus that we targeted for analysis, as well as the key elements of the text-mining approach we implemented (more technical details are available in the online *SI Appendix*). We then present the results of the topic-modeling analysis, with an overview of the topics we discovered in the corpus and that we grouped into broader categories. This first set of results is complemented by diachronic analyses that reveal how topics evolved in significance over the last three decades in the published works. We then discuss advantages and limitations of the methodology. Finally, we put the results in perspective with previous studies, from the point of view both of the topical content of the philosophy of biology and of the evolution of this content over time.

Corpus and methodology

Text-mining approaches—including topic-modeling algorithms—exploit the fact that words are not used at random in texts in which their authors attempt to convey meaning. Because words are used in specific combinations with one another, they tend to form repeated patterns wherever they occur in texts. Hence the intuition that studying these patterns can provide insights into the semantic content of these texts. Though only targeting texts at the superficial level of their lexicon, text-mining approaches aim at revealing deeper underlying semantic regularities. Such approaches are usually justified by appealing to the coherent use of words by authors, and by the marginal presence of ambiguous and metaphoric language in the scientific discourse (Firth 1957). Algorithmic text-mining approaches have been devised precisely to analyze word patterns in digital text corpora (e.g., Srivastava and Sahami 2009; Aggarwal and Zhai 2012). In particular, topic-modeling algorithms identify sets of words that occur with similar associative patterns in a given corpus and group them into topics. In turn, analysis of these topics—most notably of their most significant words and most closely related documents—make it possible to investigate the thematic content of that corpus. By considering publication years, diachronic analyses can also be carried out to assess the evolution of topics and their significance over time. It is this overall methodological approach that we implemented to identify the topics of *Biology and Philosophy* and their evolution in the last three decades. The main algorithm we used is based on the well-known Latent Dirichlet Allocation (LDA) model, which is part of a larger family of unsupervised statistical algorithms for topic discovery in texts (Pritchard et al. 2000; Blei et al. 2003). Such algorithms make it possible to explore corpora without any a priori content-related knowledge, and notably without any pre-conceived idea of which topics might be present or not. In this context, each topic is a probability distribution over every word type that is present in the corpus, and, in turn, each document is characterized by a probability distribution over topics. One way to make sense of this approach is to imagine "constructing" an article first by drawing a topic (on the basis of the document probability distribution over topics), and then by drawing a word from the topic (on the basis of topic probability distribution over words). Through an iterative convergent process, LDA algorithms assess the best probability distributions of words within topics and of topics within documents, thereby making it possible to retrieve the underlying "latent" patterns that words and topics form within the corpus' set of documents. The topic-modeling method we used can be described in five main stages.

(1) *Corpus retrieval and cleaning* We retrieved all journal articles of *Biology and Philosophy* in full-text and in electronic format from the publisher platform—Springer—from 1986 when it was launched up until 2017, hence a total of 32 complete years. Removing editorials, announcements, book reviews and other notices from the 1511 downloaded documents resulted in a corpus of 1060 articles (including 121 discussion articles of length comparable to that of regular articles). The total number of words for these 1060 articles amounted to over 8 million words (more exactly: 8,331,560 words). As can be seen in Fig. 1, discussion articles were typical of the journal issues in the 1980s and early 1990s while regular articles increased in frequency throughout the same period.[2] Also, as time went by, not only were there more articles published per year, but their length increased as well.[3]

(2) *Data preprocessing* This stage consisted in preparing the corpus in a suitable form for computational analysis. Data preprocessing included a lemmatization-based spelling normalization step during which the spelling of words was homogenized.[4] In addition, because not all types of words are proper candidates for expressing topics and may introduce noise (for instance: determinants, prepositions or pronouns, as well as rare words), we filtered them out. We also removed citations wherever we could (identified by proper names and dates). All these operations were done with

[2] Discussion articles in the early years of the journal were seen as a means to introduce people to the field and get them engaged (Ruse, personal communication).

[3] The number of yearly articles went from 20 to 30 articles per year in the first 10 years of *Biology and Philosophy* to some 30–50 articles per year in the most recent decade (over the 32 years we examined, the average amounts to 35.3 articles per year). Meanwhile the average number of words per article increased from 7–8000 in the second decade to 8–9000 in the last decade.

[4] More specifically, lemmatization is the process of grouping together the inflected forms of a word based on their intended meaning so they can be analyzed as a single item, identified by the word's lemma, or dictionary form (for instance, the lemma for "evolved" is "evolve").

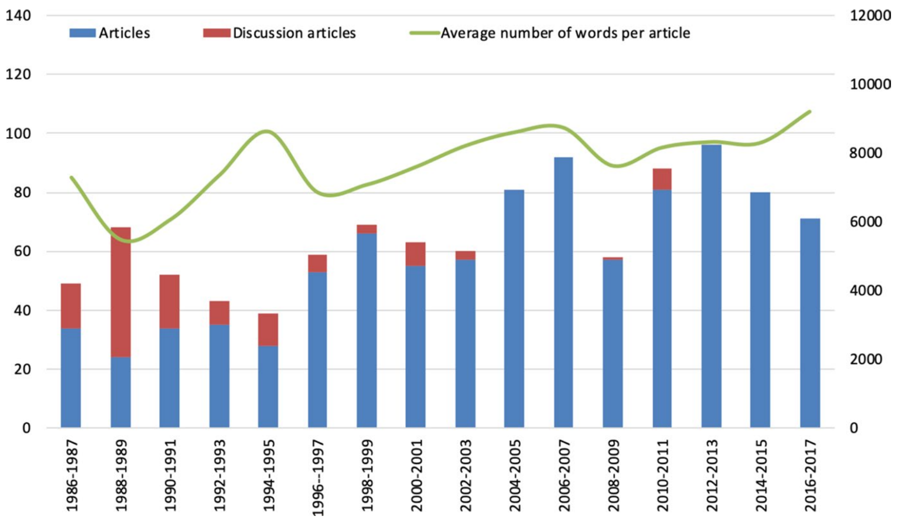

Fig. 1 Number of *Biology and Philosophy* articles published per 2-year period (left axis, regular articles in blue and discussion articles in red) and average number of words per article (right axis) (total number of documents: 1060)

the help of specific algorithms (See *SI Appendix*). This resulted in a lexicon of 9588 distinct words.

(3) *Topic modeling* This stage consisted in implementing a well-known Latent Dirichlet Allocation (LDA) algorithm, following (Blei et al. 2003). As mentioned earlier, LDA is a generative probabilistic computational method that models topics as probability distributions over words, and documents as probability distributions over topics in a corpus. The approach assumes that a hidden "latent" set of topics exists in any given corpus, and that such a set of topics accounts for the distribution of words and their cooccurrence patterns in the documents of that corpus. Through an iterative probabilistic approach, LDA makes it possible to infer sets of topics that best fit these cooccurrence patterns. Starting from random probability distributions and an assumed number K of topics, the topic modeling adjusts the probability distributions (i.e., statistically learns) up until a convergence criterion is met. The LDA—like all other well-established topic-modeling approaches—assumes that the dimensionality K of the topic distribution is known beforehand. Though heuristics exist to narrow down the range of possible values, the proper number K of topics is ultimately assessed by trial-and-error. In the present case, after comparing the resulting topics of several models (50, 75, 100, 150), we settled on $K=75$ topics. The topic-modeling stage thereby resulted in 75 probability distributions over the corpus lexicon (probabilities of finding specific words in any specific topic), and of 1060 probability distributions over the 75 topics (probabilities of finding specific topics in any specific article).

(4) *Topic interpretation and categorization* Formally speaking, in LDA, topics are probability distributions over the lexicon words of the corpus. These distributions are called "topics" because one can interpret these distributions as such: looking at

the most probable words, one can usually infer the semantic content that these words are supposed to convey. To confirm or infirm such inferences, one can retrieve the documents in which specific topics are the most probable. Though manual and error-prone, interpreting topics and assigning them labels is strongly constrained by their sets of words and the documents in which they appear. We interpreted and labeled all 75 topics with the same approach, by looking both at their respective sets of words and at the most relevant documents for each topic. Note that we tried to stick, as much as possible, to labels that corresponded to the top-words for each topic. In a second step, we grouped topics with related content into categories. This grouping was done manually on the basis of the topic interpretation and our knowledge of the field. We also used a topic correlation matrix as a heuristics to probe the validity of our groupings (See *SI Appendix*).[5] Among the 75 topics, we identified 8 that we categorized as "jargon": these topics appeared to be either too generic or polysemic to be precisely related to any meaningful philosophy of biology issue.[6] The presence of such jargon topics was to be expected since topic-modeling algorithms run over all words present in the full-text corpus, including those words that are used to package ideas into meaningful sentences. In fact, these groups of words are even quite significant, with a cumulated probability of being found in the corpus that amounts to some 0.37 (see figure S2 in the *SI Appendix*); yet, because these topics are not specific to any well-identified research theme and tend to be spread out all over the corpus, we have set them aside in what follows. For the rest of our analysis, we focused on the remaining 67 topics, grouped into a manageable number of 9 categories.

(5) *Diachronic topic analysis* The final stage of the method consisted in analyzing the diachronic distribution of topics over the 32 years of *Biology and Philosophy*, from 1986 to 2017. In order to average out possible year-to-year variance, we chose to split the corpus into 16 periods of 2-years each. For every topic, we computed its probability of being found in any given period simply by averaging the probability of finding that topic in all articles of the period. Then, by averaging over the probabilities per time-period of all topics belonging to any given category, we also computed the probability of finding that category at any given time-period. In this way therefore, we were able to quantify the relative frequency of topics and of categories at any given time-period, hence our assessment of their temporal evolution.

[5] These categories are therefore not an algorithmic outcome of the topic-modeling methods, but result from our best judgement about how topics relate to one another. The reason for proposing such a grouping of topics into categories is very pragmatic: it is a means to handle the overall high number of topics. Of course, topics relate to one another in multidimensional ways, depending on the relative probabilities of the words that best express them. There are therefore multiple ways of grouping topics into categories and other choices than the ones we made are also possible. Yet, one shone should simply bear in mind that categories are made for convenience: in topic-modeling analyses, only topics ultimately matter.

[6] For instance, topic #9 Way-Think-Make included such words as "think; make; say; good; case; just; like; need; kind; thing", topic #42 Theory-Argument-Claim included "theory; argument; claim; question; view; term; fact; point; follow; principle; sense" (See Table S1 in the *SI Appendix* for more details). Though of no direct use for the purpose of the present analysis, these jargon topics could however provide interesting insights on the generic features of the philosophical discourse and style of argumentation.

The topics of *Biology and Philosophy*

The first obvious result of our analyses is that the topics found in *Biology and Philosophy* are extremely diverse: they cover a broad range of interests, ranging from evolution—as could be expected…— and systematics to cognition, socio-normative issues and general philosophy of science questions (see Table 1). This diversity of topics is also visible on the topic-correlation network of Fig. 2.

Looking first at the relative frequency of categories—as measured by the sum of their topic occurrence probabilities—the two most significant categories are those that concern *general philosophy of science* (category *H,* at 22%) and *evolution* (category *A*, at 18%), with a combined occurrence weight in the corpus of 40% (see Fig. 3, once the category *J-Jargon* has been set aside). They are also the two categories with the highest numbers of topics (respectively 14 and 11, as can be seen in Table 1). All other 7 categories are rather balanced in terms of probability in the corpus (between 6 and 10%) and of numbers of topics (from 5 to 7).

More specifically, we grouped in category *A-Evolution* topics that directly relate to natural selection and evolution, including evolutionary processes, patterns, the role of contingency (topics #1 Natural Selection, #33 Darwin-Selection, #34 Evolutionary-Change), but also more specifically to the concept of fitness and its relevance to organisms and populations (topics #54 Fitness-Individual-Group, #73 Fitness-Organism, #67 Fitness-Population). The category also includes topics that concern the notions of adaptation and extended phenotype, as well as that of replicators, variation and inheritance (topics #15 Constraint-Adaptation, #56 Environment-Adaptation, #31 Extended Phenotype, #62 Replicators). Finally, also included in this category is a topic about drift (topic #61 Drift).

The second category, *B-Individuality-Altruism*, also includes topics that are relevant to the debate over evolution by natural selection, yet because these topics offer distinctive perspectives, we chose to put them in a separate category. In particular, one finds topics that concern evolutionary game theory, cooperation and altruism (topics #52 Altruism-group-selection, #65 Cooperation and #49 Evolutionary game theory), as well as levels of selection, organisms, groups, populations (topics #45 Individual-organism and #11 Individual-population).

Category *C-Species-ecology* covers topics that specifically deal with the concept of species (topic # 17 Species) and also by extension taxonomy and phylogenetics (#27 Character-taxon and #26 Tree-lineage). Because biodiversity-related questions often concern species—and thereby tend to mobilize, in their articles, some topics that they share with articles about systematics—we have included, in this same category, topics about biodiversity and conservation, as well as ecology as a matter of consequence (topics #5 Biodiversity and #10 Ecology).

Topics about genetics, but also molecular biology and mendelian genetics (topics #12 Genes-DNA, #43 Genes-traits, #53 Proteins and #64 Mendelian genetics) have been grouped into category *D-Genetics-Development*, alongside with topics that are more development-related (topics #47 Development-Evolution,

Table 1 List of all 75 topics sorted by category

Categories	Topics (with topic ID)	Number of topics per category (total 75)
A-Evolution	Constraint-adaptation (15); Drift (61); Environment-adaptation (56); Darwin-selection (33); Evolutionary-change (34); Extended phenotype (31); Fitness-individual-group (54); Fitness-organism (73); Fitness-population (67); Natural selection (1); Replicators (62)	11
B-Individuality-Altruism	Altruism-group-selection (52); Cooperation (65); Evolutionary game theory (49); Individual-organism (45); Individual-population (11)	6
C-Species-Ecology	Biodiversity (5); Character-taxon (27); Ecology (10); Species (17); Tree-lineage (26)	5
D-Genetics-Development	Development-Evolution (47); Development-Innateness (68); Development-Morphology (8); Genes-DNA (12); Genes-traits (43); Mendelian genetics (64); Proteins (53)	7
E-Network-entropy-information	Complexity-entropy (22); Control-network-regulation (3); Control-structure (21); Information (51); Signaling (74)	5
F-Cognition-behavior	Behavior (28); Behavior-sex (37); Beliefs (48); Cognition (7); Language (71); Mental states (69); Neurosciences (25)	7
G-Socio-normative issues	Cultural evolution (20); Economic-choice (72); Evolutionary ethics (44); Morality (60); Norm-rationality (30); Sociality (39)	5
H-General philosophy of science	Causation (46); Diagram-representation (6); Explanation-evolution (23); Explanation-mechanism (24); Hypothesis-evidence (18); Laws of nature (58); Model-equations-dynamics (55); Model-optimality (14); Probability (40); Property-kind (50); Reductionism (19); Science (59); Scientific-theory-discovery (0); Theory-model (57)	14
I-Others	Art-Modern synthesis (70); Cell biology (16); Disease-clinical (13); Function (38); Immunity-self (4); Life (63); Nature-philosophy (2)	7
J-Jargon	Approach-important (36); Case-condition (29); Increase-number-size (41); Research-biology (66); Theory-argument-claim (42); Time-change-example (32); Way-think-make (9); Work-book-science (35)	8

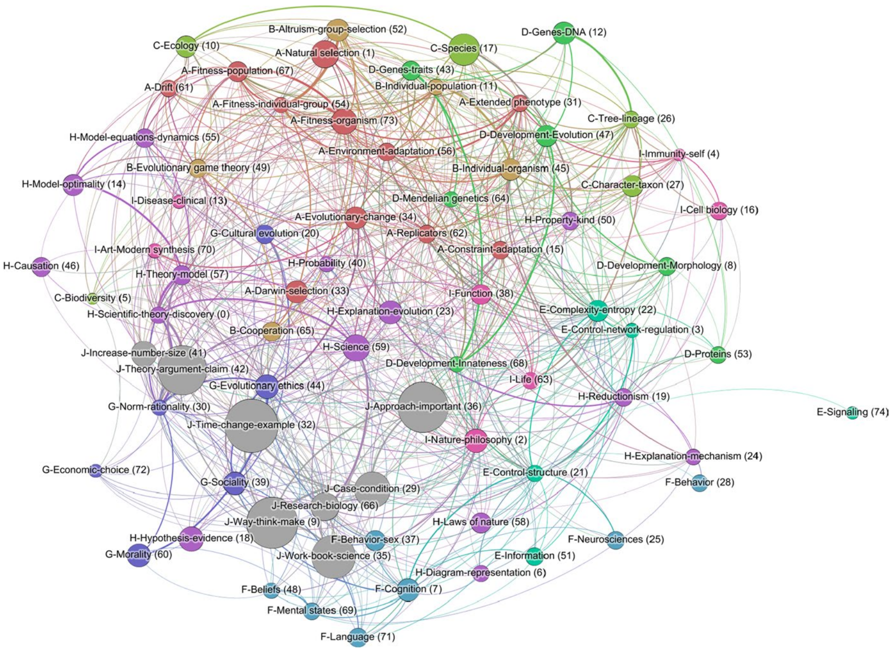

Fig. 2 Graph of the 75 topics based on their relative frequency and their correlation [nodes represent topics that have been colored based on their category; node areas are proportional to topic probabilities in the corpus; thickness of edges represent topic correlation in the corpus, with values above a threshold set to: average value plus one standard deviation; visualization tool: Gephi (Bastian et al. 2009)]

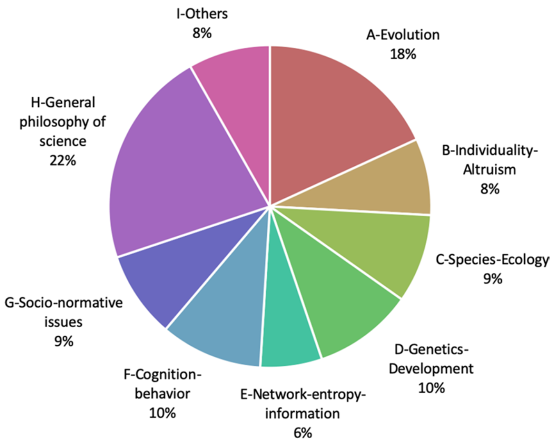

Fig. 3 Relative frequency of topic categories inside the corpus of *Biology and Philosophy* (excluding category *J-Jargon*)

#68 Development-Innateness and #8 Development-Morphology). The reason behind this grouping is that genetics-related words also appeared in development-related topics, most notably topic #47, and that phenotype-related words also appeared in topic #64.

Category *E-Network-entropy-information* is somehow less biological than the previous ones: it includes topics that relate to thermodynamics, energy, entropy, as well as information, communication and signaling (topics #22 Complexity-entropy, #51 Information, #74 Signaling). Because some of these topics included terms that related to complexity and organization, we also placed in this category topics about network, regulation and control (topics #3 Control-network-regulation and #21 Control-structure). Note that some of these topics are also related, by their sets of top-words, to topics of category *F-Cognition-behavior*.

Category *F-Cognition-behavior* includes topics that focus on the neurosciences, cognition, psychology, mental states and beliefs (topics #69 Mental states, #25 Neurosciences, #48 Beliefs, #7 Cognition) as well as topics that more specifically concern behavior and language (topics #28 Behavior, #37 Behavior-sex, #71 Language). The reason for grouping these topics together is that topics that referred to mental states also tended to include action-related words, hence behavior.

Topics of a more socio-normative content have been placed in category *G-Socio-normative issues*. These include topics about evolutionary ethics and moral judgment (topics #44 Evolutionary ethics, #60 Morality), as well as a closely related topic about norms and rationality (topic #30 Norm-rationality). We also included topics about culture and cultural evolution (topic #20 Cultural evolution) that appeared to fit better here than in category *A-Evolution*, and by extension topics about sociality and economics (topics #39 Sociality, #72 Economic-choice).

Our largest category includes more transversal topics that we interpreted as *general philosophy of science* topics (category *H*). These include emblematic topics related to causation, natural kinds, laws of nature and biological generalizations (topics #46 Causation, #50 Property-kind, #58 Laws of nature) as well as topics that concern debates about reductionism (topic #19 Reductionism), explanation (topics #23 Explanation-evolution, #24 Explanation-mechanism), as well as theory discovery and models (topics #0 Scientific-theory-discovery, #57 Theory-model, #55 Model-equations-dynamics, #14 Model-optimality). Other topics in that category concern knowledge and epistemology more broadly construed, including evolutionary epistemology (topic #59 Science); as well as knowledge- and theory-justification and probability (topics #18 Hypothesis-evidence, #40 Probability). One last topic concerns issues about representation, diagrams and the use of visual images (topic #6 Diagram-representation).

We grouped under category *I-others* a number of topics that we could not easily assign to any of the other categories. These include two topics that relate in large part to philosophy of medicine (topics #13 Disease-clinical, #4 Immunity-self) and a topic about cells and tissues (topic #16 Cell biology) which is present in articles that concern part-whole relationships in organisms, cellular organization, but also cancer and cellular evolution. Note that these last three topics are not very sharp in their thematic delineation, a sign of the limitation of the topic-modeling approach past a certain level of detail. For instance, the topic Disease-clinical (13) includes

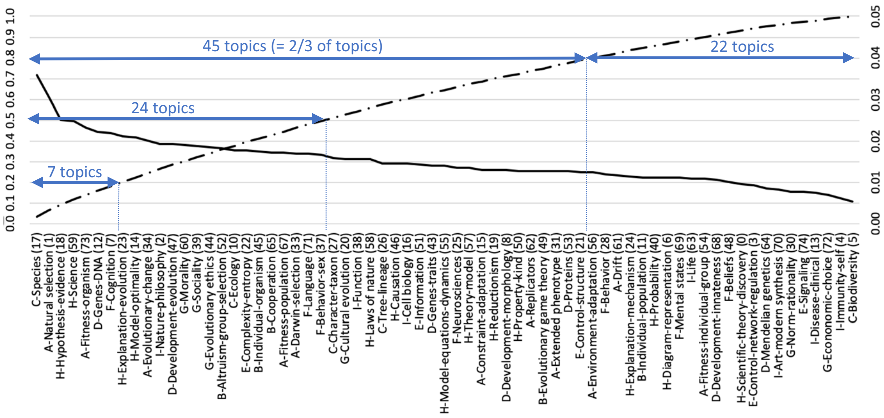

Fig. 4 Topics sorted by decreasing order of frequency (dashed line/left axis: cumulated probability of occurrence of topics in the corpus; solid line/right axis: probability of occurrence of topics in the corpus; labels include: category-letter/topic-label/topic-ID; total number of topics: 67, excluding jargon topics)

terms that refer to evidence-based medicine but also to Kettlewell and his peppered moth studies, likely linked by the role played by evidence (yet still quite distinct). Similarly, the topic Immunity-self (4) can be found in a diverse range of papers concerning immunology as well as individuality (see Tables S1, S2). Also included in this category *I-others* are topics about the notions of function and life (topics #38 Function, #63 Life), as well as a rather heteroclite topic that groups terms related to art and aesthetics present in articles about evolutionary aesthetics and musicality together with terms that can be taken to relate to the modern synthesis (topic #70 Art-Modern synthesis).[7] Finally, one finds a topic that relate mostly to papers of significant historical dimension (topic #2 Nature-philosophy). Indeed, this quite broad topic about nature, human and animal forms, objects and bodies, appears to be clearly associated with articles of a distinctive historical perspective (see Tables S1, S2).

When comparing topics in terms of their relative frequency in the corpus, it is interesting to note that the first 7 topics already represent 20% of cumulated probability (excluding jargon-related topics) (see Fig. 4). These most significant topics include topic #17 Species, two topics related to evolution (topics #1 Natural selection and #73 Fitness-organism), the topics Genes-DNA (12) and Cognition (7) as well as two general philosophy of science topics (topics #18 Hypothesis-evidence and #59 Science which is, as mentioned above, mostly about evolutionary epistemology). To reach 50% of cumulated probability, 17 additional topics are needed, hence 24 topics in total, representing all categories except category *I-Others*. At the

[7] We hesitated classifying topic #70 as a jargon topic, yet, when reviewing its most strongly associated articles, one finds entries that concern evolutionary aesthetics as well as the modern synthesis. Hence our decision to leave it in category I-Others.

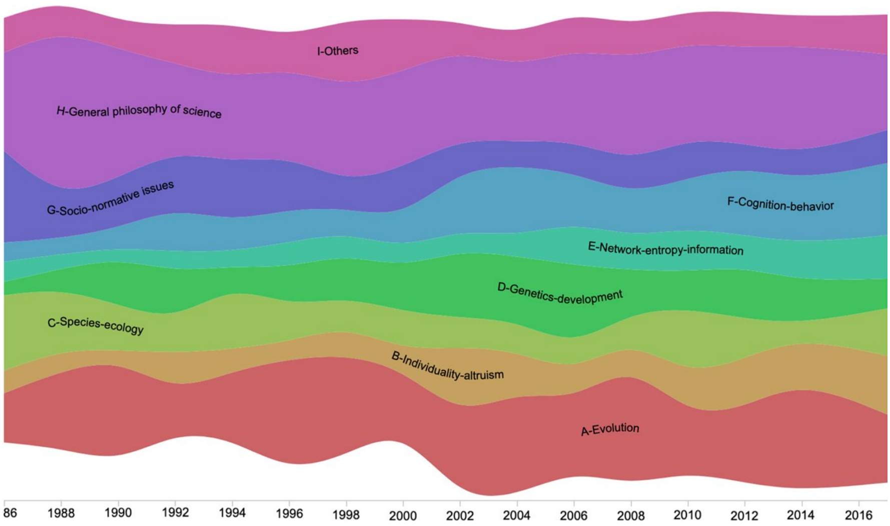

Fig. 5 Diachronic evolution of the topic-categories in *Biology and Philosophy* between 1986 and 2017 [the width of each 'stream' is proportional to the probability of each category in the corpus, jargon-topics excluded; visualization tool: RAWGraphs (Mauri et al. 2017)]

other end of the spectrum, topics that contribute to the last 20% of cumulated probability are more numerous and include 22 topics. Among the least probable topics, one finds Biodiversity (5), Immunity-self (4), Economic-choice (72) and Disease-clinical (74). All in all, about two-thirds of the topics contribute 80% of cumulated probability, while the remaining third only contributes 20%, and within these two-thirds of topics, all 9 categories are well represented.

How did *Biology and Philosophy* change from 1986 until today?

Research themes evolve over time, and so does *Biology and Philosophy*. A journal is also subject to changes imparted by editorial policies and the broader publishing context within which it develops, including publisher strategies and competitor journals. Independently of the sources of change, text-mining methods make it possible to identify the evolution in topic significance over time. One can thereby identify which topics were prevalent and when. One can also analyze diachronic trends at the aggregate level of topic categories. Because topics correspond to the sets of words used by philosophers in their papers at different time-periods, they also reflect the types of research questions that received attention and got published during these periods. A broad overview of the evolution of topics is depicted in Fig. 5, which represents the probability of finding corresponding categories in *Biology and Philosophy* at specific time-periods. Even at such high level of granularity, the results show interesting diachronic patterns, notably the fact that not everything in *Biology and Philosophy* is always about evolution. It is true that evolution-related topics broadly construed—that is to say, including topics of categories *A-Evolution,*

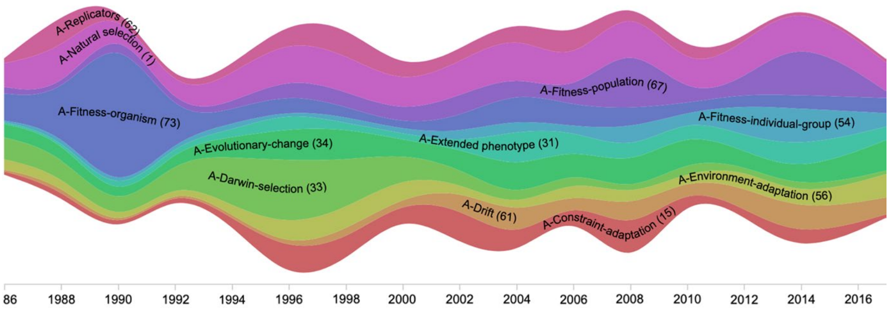

Fig. 6 Diachronic evolution of the topics of category *A-Evolution* between 1986 and 2017 (the width of each 'stream' is proportional to the probability of each topic in the corpus)

B-Individuality-Altruism, C-Species-Ecology and *D-Genetics-Development*—have always represented the lion's share of the journal publications. Yet it is also interesting to note that the relative significance of these categories has encountered ups and downs in the past 30 years, notably with a stronger preponderance in the early 2000s than in the mid 2010s. Two categories appear to have more regular and slightly upward trends: these are the categories *E-Network-entropy-information*, *F-Cognition-behavior*, which both tend to increase over time. On the other hand, category *G-Socio-normative issues* exhibits an overall decreasing pattern. Category *H-General philosophy of science* tends to be relatively constant over time, after having decreased from an early burst in the 1980s. Finally, category *I-Others* appears to be relatively constant overall, despite a slight swelling in late 1990s.

If one zooms in onto category *A-Evolution* (Fig. 6), one can visualize temporal changes at the level of topics that reveal a relative constancy of such topics as Natural-selection (11) or Environment-adaptation (56), but also more contrasted patterns. For instance, Fitness-organism (73) appears to have been much discussed between 1988 and 1992, but much less so afterwards, being somehow replaced by discussions of fitness at the group or population levels around 2008 and 2014 (topics #67 Fitness-population and #54 Fitness-individual-group) and an interest in extending the basis for fitness (topic #31 Extended phenotype).[8] One also notes a slight disinterest from the replicators view since 2012 (topic #62 Replicators) and from studying selection in relationship to Darwin himself or his views since the late 2000s (topic #33 Darwin-selection). In parallel, discussions about drift, which started in the late 1990s, have been relatively steady since then (topic #61 Drift).

Results for topics of categories *B-Individuality-Altruism* and *C-Species-Ecology* are shown on Fig. 7. Of particular salience are several burst-like patterns, notably for topics that relate to the species problem in the late 1980s and around 1994 (topic

[8] Note that some of the bursts may be due to special issues. For instance, issue 2 of volume 29 (2014) was a special issue on formal Darwinism, while issue 5 of volume 23 (2008) was a special issue on the adaptive landscape. This shows in the diachronic evolution of the topic Fitness-population (67) which displays peaks in 2008 and 2014. Similarly, issue 1 of volume 6 (1991) included papers and invited commentaries on the notion of fitness; hence (part of) the peak of topic Fitness-organism (73) around 1990.

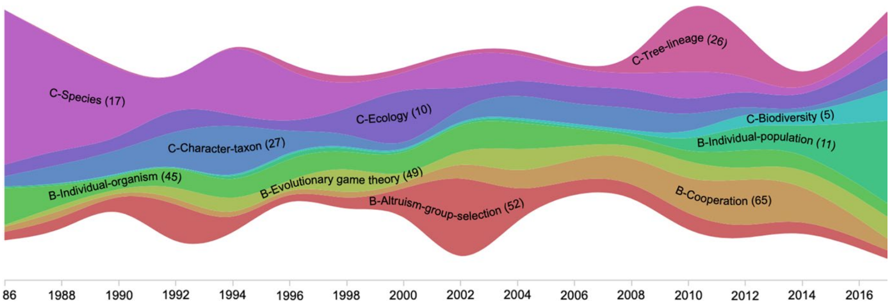

Fig. 7 Diachronic evolution of the topics of categories *B-Individuality-Altruism* and *C-Species-Ecology* between 1986 and 2017 (the width of each 'stream' is proportional to the probability of each topic in the corpus)

#17 Species), and to taxonomy also around 1994 and phylogenetics in the early 2010s (topics #27 Character-taxon and #26 Tree-lineage, the latter being strongly correlated with a 2010 special issue on the tree of life). Publications about altruism and group selection appear to have peaked around 2002, and been somehow replaced by an interest in cooperation (topics #52 Altruism-group-selection and #65 Cooperation). Discussions about individuals versus populations have tended to increase lately (notably topic #11 Individual-Population). A similar recent interest in biodiversity also seems to have occurred (topic #5 Biodiversity), possibly catching up from a burst of interest in ecology around 2000 (topic #10 Ecology).[9]

Figure 8 shows the temporal evolution of the topics of the categories *D-Genetics-development* and *E-Network-entropy-information.* While topics related to genetics broadly construed (topics #64 Mendelian genetics, #12 Genes-DNA, #43 Genes-traits and #53 Proteins) have received a relatively stable interest, especially since the 2000s, a burst of interest in development-related topics is clearly visible between 2000 and 2006, notably linked to articles on evo-devo questions (topics #47 Development-Evolution, #68 Development-innateness and #8 Development-morphology). One can also note a fairly recent increase in information-related topics, including signaling and control since 2006 (topics #51 Information, #74 Signaling and #21 Control-structure).

Two opposite trends appear to characterize the categories *F-Cognition-behavior* and *G-Socio-normative issues* (Fig. 9): while the former has seen a strong rise for the last 30 years, the latter has tended to decline in relative significance. This can be seen at the topic level. Topics that relate to behavior broadly construed (topics #78 Behavior and #37 Behavior-sex) or to cognition, including language and the neurosciences (topics #48 Beliefs, #7 Cognition, #71 Language, #69 Mental states and #25 Neurosciences) have all generally increased over the past three decades. Conversely, topics about evolutionary ethics and morality have tended to decrease over

[9] Several ecology-related articles appeared in the 2000 issue 16(4), though there is no apparent indication of that issue being a special issue on the topic.

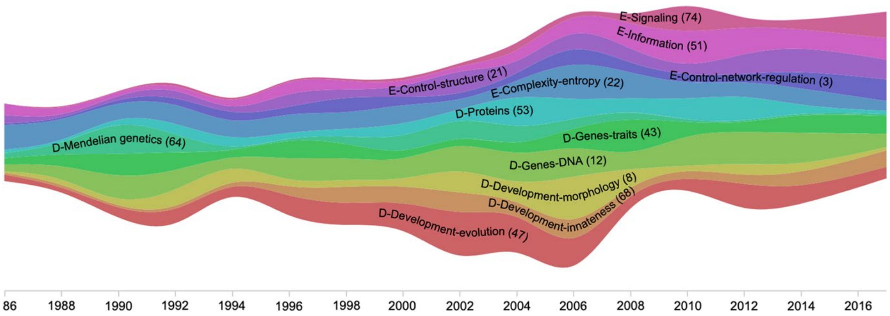

Fig. 8 Diachronic evolution of the topics of categories *D-Genetics-development* and *E-Network-entropy-information* between 1986 and 2017 (the width of each 'stream' is proportional to the probability of each topic in the corpus)

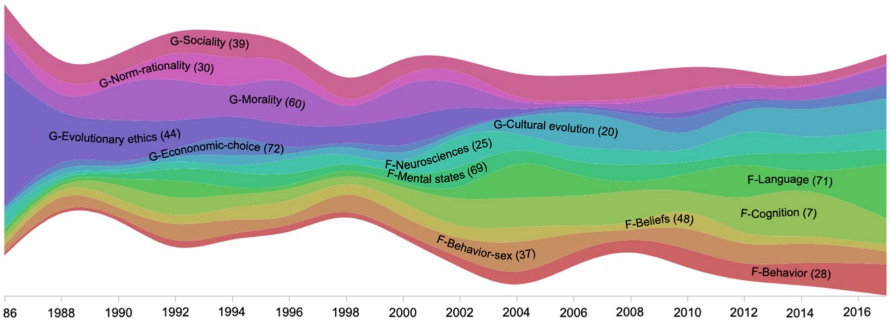

Fig. 9 Diachronic evolution of the topics of categories *F-Cognition-behavior* and *G-Socio-normative issues* between 1986 and 2017 (the width of each 'stream' is proportional to the probability of each topic in the corpus)

the same time-period (topics #44 Evolutionary ethics and #60 Morality) and have been barely compensated by an increase in interest in cultural evolution from 2004 onward (topic #2 Cultural-evolution).

Besides a slight prominence around 1988, category *H-General philosophy of science* has remained fairly constant in significance (Fig. 10). However, individual topics have changed. One of the most noticeable changes concerns the marked decline of topic Science (59), which relates to evolutionary epistemology and scientific knowledge, including scientific change.[10] Other topics that have declined in the past three decades include topics about scientific discovery and about reductionism (topics #0 Scientific-theory-discovery and #19 Reductionism). On the other hand, model-related topics have tended to increase in significance over the years, especially since the 2000s (topics #55 Model-equations-dynamics, #14 Model-optimality). Also, when looking at explanation-related topics,

[10] It is likely that the peak observed around 1988 (partially) results from a paper by David Hull on evolutionary epistemology and invited commentaries, all published in issue 3(2) of 1988.

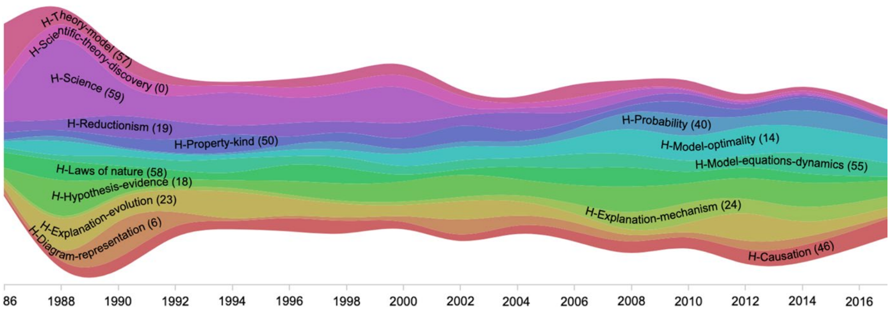

Fig. 10 Diachronic evolution of the topics of category *H-General philosophy of science* between 1986 and 2017 (the width of each 'stream' is proportional to the probability of each topic in the corpus)

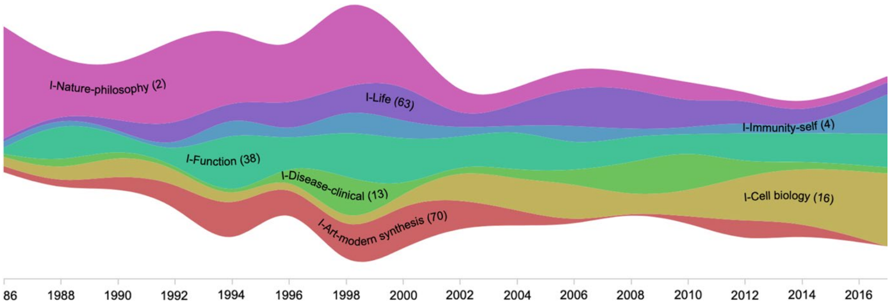

Fig. 11 Diachronic evolution of the topics of categories *I-Others* between 1986 and 2017 (the width of each 'stream' is proportional to the probability of each topic in the corpus)

mechanistic explanation can clearly be seen as emerging in 2000 and developing onward (topic #24 Mechanism-explanation), in interplay with evolution-related explanation, notably on questions about the distinction between ultimate and proximate causes (topic #23 Explanation-evolution).

In category *I-Others*, one immediately notices how the topic Nature-philosophy (2) decreased in frequency from the 2000s (Fig. 11). Remember that this quite broad topic is strongly associated with articles with a distinctive historical dimension. The evolution of this topic thereby reflects the frequency of historically-minded papers in *Biology and Philosophy* between 1986 and the 2000s, and their strong decline since then. The topic Function (38)—which clearly depicts research themes on the notion of function and its different accounts—appears to have remained of a somehow constant interest, after slight bursts in the 1990s. A topic that has seen an increase since the 2000s is the topic Cell biology (16), which is present in articles that concern part-whole relationships in organisms, cellular organization, but also cancer and cellular evolution. Two topics related to philosophy of medicine issues (topics #13 Disease-clinical and #4 Immunity-self) have witnessed fluctuating significance, though, as we noted earlier, the conceptual delineation of these topics is not very good. Something similar can be said about the topic Art-modern-synthesis (70), which groups terms that are present in articles about the modern synthesis, especially throughout the 1990s, and terms present in articles about evolutionary aesthetics and musicality around 2012 (see Table S2). Finally, the topic Life (63), which is found both in design- and creationism-related articles and in origins of life articles, has seen two slight peaks, first in the 1990s and then the late 2000s.

As could be expected, the diachronic analyses at the level of individual topic are more contrasted than at the aggregate level of categories. Upwards and downward trends are clearly noticeable, and so are the interests of philosophers of biology over the last three decades.

Methodological advantages and limitations

Topic-modeling is a well proven approach to identifying topics from corpora of digital texts, most notably with the assistance of an LDA-based algorithm (Pritchard et al. 2000; Blei et al. 2003; Griffiths and Steyvers 2004; Griffiths et al. 2007; Blei and Lafferty 2009; DiMaggio et al. 2013). Such methods have two notable advantages. First, they make it possible to carry out semantic investigations over large corpora whose analysis would have been otherwise impossible by hand. Second, they work in a non-supervised data-driven way, meaning that topics are identified bottom-up on the basis of the content of the corpus and without a priori knowledge of which topics populate the corpus.

This doesn't mean, though, that everything happens hands-free: there are technical choices to be made at all stages of the methodology, including iterative rounds of parameter settings, simulations and result inspection (Hu et al. 2014). For instance, in the first stage of corpus cleaning, automatic filters on document metadata were complemented by manual inspections of the corpus so as to select which documents to filter out (e.g. editorials, errata etc.). For data preprocessing (stage 2), this included inspecting the lexicon that resulted from lemmatization processing and word filtering (visual inspection aided by specific queries for stop-words, special characters and word frequencies). For the topic modeling stage (stage 3), as mentioned earlier, we ran analyses for different values of the number of topics K (50, 75, 100, 150), before settling on $K=75$. That value appeared to us as the optimum to avoid redundancies between topics (more frequent for high values of K) and semantic fuzziness of topics (more frequent for low values of K).[11] As explained above, stage 4 of topic interpretation and categorization is a highly manual stage which involves semantic inferences and choices on the basis of word sets and article examples. Finally, the diachronic topic analyses of

[11] Though with 75 topics, we still had to deal with a few topics that would have benefited from being merged and a couple others that would have been better split. Yet we found that granularity to be the best trade-off. Note also that we found a strong robustness in the results from one value of K to another: many topics from one model can be made to correspond to topics or groups of topics from another, thereby lending confidence in the topic-model that was ultimately chosen.

stage 5 also involved choices in terms of periodization (we chose 2-year periods). Topic modeling results should therefore always be relativized to these technical choices. They should also be relativized to the corpus that has been chosen. In the present case, the study only concerns one single journal. Inferring general trends about the philosophy of biology in general from the analyses we conducted solely on *Biology and Philosophy* should therefore be done with caution, all the more so as philosophers of biology also publish in other journals (general philosophy of science journals, general philosophy journals and—increasingly—biology journals), under other forms than articles (e.g. monographies, edited volumes), and in different languages (other than English). Having said that, *Biology and Philosophy* is arguably one of the most significant journals in the field, and the analyses we conducted certainly do bring insights into the journal intellectual content and its evolution over time.

As we saw earlier, topic-modeling approaches provide detailed viewpoints both from synchronic and diachronic perspectives. One should always remember that the methodology only results in identifying topics (as probability distributions over the lexicon terms, based on for their patterns of co-occurrence in the documents) and their probability of occurrence in the documents of the corpus. One of the limitations of the methodology is that it is not meant to capture more sophisticated relationships between words or even between topics, such as entailment or causal relations [other approaches should be implemented, such as conceptual analysis or argument mining (Peldszus and Stede 2013; Swanson et al. 2015)]. Therefore, while the methodology makes possible certain analyses in terms of topic cooccurrences in documents or topic evolution, it cannot reveal deeper relationships between topics that authors intended in their papers, nor the filiation of topics over time.

That being said, topic-modeling approaches provide quantitative perspectives—for instance on the development of a discipline—that are complementary to classical historical methods. One of their most significant benefits is the non-supervised way in which they function, thereby providing a fairly objective vantage point onto the topics of the philosophy of biology and their historical evolution as seen in the corpus of *Biology and Philosophy*. It is of course quite reassuring to find topics and associated papers that make sense to the expert (see for instance table S2) or that fit well-known facts about the discipline in general (such as the significance of evolution-related works, the spike in evo-devo interest or the appearance of the mechanistic model of explanation in the 2000s). Yet there was no guarantee at start that one would find these patterns: the topics were identified via purely data-driven approaches. This lends fallibility to the approach, but also empirical strength to the results.

The results in perspective

A plurality of topics in *Biology and Philosophy*

One of the main conclusions that Gayon (2009) drew from manually sorting the 1986–2002 articles of *Biology and Philosophy* was the predominance of evolution-related works which accounted for some 35% of all articles. Pradeu (2017) inferred even stronger results, assessing that between 55 and 62% of articles—respectively over the periods 1986–2002 and 2003–2015—concerned evolution broadly construed, including systematics and human evolution. Our own analyses in terms of topic coverage show that the share of evolution-related topics among all topics of *Biology and Philosophy* appear to be roughly closer to some 35% when one cumulates the relative frequency of all topics of categories *A-Evolution* (18%), *B-Individuality-Altruism* (8%), and *C-Species-Ecology* (9%).[12] Of course, the two types of analyses are hard to compare since, in the cases of Gayon and Pradeu, numbers represented percentage of articles, whereas in our case they capture the probabilities of finding topics within all articles of the corpus. Nevertheless, our results show that the predominance of evolution as a research topic is roughly in the order of 1/3 of all research topics, and that there is much more to philosophy of biology than philosophy of evolutionary biology. This is something that Hull was clearly aware of: despite recognizing the prevalence of selection and the species problem, he was also cognizant that the field included many more topics, including evolutionary ethics and evolutionary epistemology, ecology, genes and development, not to mention more specialized topics targeted by special issues such as models of scientific theories, individuality, functional explanation, pictorial representation, or the neurosciences (Hull 2008, 28–29). It is this diversity that our analyses also reveal.

Indeed, topic-modeling algorithms make it possible to capture details at a much finer grain, revealing not just a diversity of topics outside evolution-related categories but also within (as can be seen in Table 1 above). Furthermore, instead of being reduced to a single category, articles are modeled as probability distributions over topics, hence the possibility for any single article to exhibit several topics. We believe such a view is closer to the reality of philosophy of biology research—and more generally research in philosophy of science—that weaves together theories, concepts, case studies from the different scientific disciplines and research questions that are specific to the practice of some form of analytic philosophy. This shows in how topics can mix in many different ways in the corpus articles. For instance, among the top-20 articles for topic #1 Natural selection, one finds a typical evolution-related article by Roberta L. Millstein—her 2002 article "Are Random Drift and Natural Selection Conceptually Distinct?"—which is, as could be expected, also strongly affiliated with topic #61 Drift. Yet one also finds the 2008 article by Stuart Glennan entitled "Productivity, relevance and natural selection". That article is not only strongly related to topic #1 Natural selection but also to topic #46 Causation, and to a lesser extent #61 Drift. Such an article clearly is at the crossroads of evolutionary theory and general philosophy of science issues. If, in turn, one looks at the top-20 articles for topic #46 Causation, one finds the 1999 article of Lisa Gannett, "What's in a Cause?: The Pragmatic Dimensions of Genetic Explanations". Examining the top-topics for that article, in addition to topic #46 Causation, returns topics #12 Genes DNA and #43 Genes Traits, which are then quite removed from evolutionary biology studies. This shows that topic #46 Causation clearly is a transversal

[12] If we try to stick more closely to the categories proposed by Gayon and adapted by Pradeu, and add the four topics that concern evolutionary-epistemology and evolutionary ethics from categories *G-Socio-normative-issues* and *H-General-philosophy of science*, while removing the two ecology-related topics of category *C*, then the share of broadly construed evolution-related topics slightly increases to 38%.

topic whose significance would have most likely gone unnoticed—or at best underestimated—in single-category article classifications.

Topic-modeling approaches therefore make it possible to get a fairly detailed and nuanced view about the diversity of research topics that philosophers of biology investigate. This level of detail also shows when one compares findings for topics that concern what Gayon had labeled "philosophical questions of general interest regarding biology and the living world" (Gayon 2009). Indeed, though Gayon had grouped under this category many different items, the level of granularity of his analyses prevented him from specifying the details of what was included and what not, and in which proportions. A similar coarse-graining taxonomy of philosophy of biology enquiries was proposed by Griffiths, who distinguished three major strands of research: (1) general theses in the philosophy of science that are addressed in the context of biology, (2) conceptual puzzles within biology itself that are subjected to philosophical analysis, and (3) appeals to biology that are made in discussions of traditional philosophical questions (Griffiths 2008). While the topics we identified could be made to fit these three broad categories—for instance, topics #19 Reductionism in (1), #73 Fitness-organism in (2), or #60 Morality in (3), as Griffiths does—these categories are quite permeable (as are our own categories). This is notably the case with topics that can be assigned to more than one category, for instance topic #60 Natural selection that may be assigned to (1) if taken in the context of a discussion on the semantic view of theories, or to (3) if in the context of a questioning on evolutionary ethics (as already recognized by Griffiths). This also shows in the different assemblages of topics inside articles, which sometimes make it hard to classify individual articles in specific categories. By targeting a lower level of granularity, text-mining approaches make more details available, providing the basis for more fine-grained analyses, including also more nuanced diachronic patterns in topic interests (as we will see below).

Note that the topics—as lists of words—have been automatically generated by the topic modeling algorithms in a bottom-up way, solely on the basis of the content of the articles. As a consequence, articles are never 'forced' to fit into pre-existing categories, but rather exemplify specific sets of self-generated topics. It is therefore not surprising that our topic classification differs from those used by Gayon and by Pradeu. Yet the bottom-up approach of the methodology gives good reasons to believe that the results are closer to the actual content of *Biology and Philosophy* articles. In a way, topic-modeling is to article classification what computational phylogenetics is to taxonomy. It also shows the limits of sorting articles into a pre-defined classificatory grid.[13]

[13] This is most apparent with the *PNAS* disciplinary classification that was superimposed onto *Biology and Philosophy* articles in Pradeu (2017). This was done with a view to contrasting the focus of philosophers of biology onto evolution-related topics compared to the broad diversity of scientific publications in all biological disciplines. Yet, while *PNAS* categories do capture some of the biological disciplines that philosophers of biology are interested in, they do not do justice to the diversity of research interests that philosophers pursue. Investigating this diversity of topics also shows that one cannot expect philosophy of biology articles to reflect the relative importance of scientific publications in the different biological disciplines: whereas some of the topics that we identified can reasonably be mapped onto disciplines of biology—for instance topics such as #1 Natural selection (evolutionary biology) or #12 Genes DNA (genetics)—many others resist such mapping, such as topic #46 Causation that we discussed above as

Diachronic evolution of *Biology and Philosophy*

Another contrasting point of our studies compared to previous ones is whether *Biology and Philosophy* has changed in topical content over the last three decades or not. While our broad results—at the level of topic-categories—tend to agree with Gayon's and Pradeu's when it comes to documenting a relative stability of interest in evolution-related topics (see category *A-Evolution* in Fig. 5), they disagree in the case of other categories, which showed a different diachronic picture and did change over time, some in more pronounced ways than others (for instance, categories *B-Individuality-Altruism*, *E-Network-Entropy-Information* or *F-Cognition-Behavior* have tended to increase in frequency, while categories *C-Species-Ecology* or *G-Socio-Normative-Issues* have tended to decrease). But the changes are even more visible at the level of individual topics as we noted above (see Figs. 6, 7, 8, 9, 10, 11). Interest in topics such as Species (17), Mendelian Genetics (64), Evolutionary ethics (44), as well as papers about evolutionary epistemology (topic #59 Science) or with a more distinct historical dimension (topic #2 Nature-Philosophy) have clearly tended to decrease. At the same time, topics such as Cultural evolution (20), Behavior (28), Cognition (7) or Cell biology (16) received increase attention, while other topics had fluctuating destinies with more or less pronounced ups and downs (e.g. topics #52 Altruism-group-selection, #8 Development-morphology, #71 Language or #38 Function).[14]

One possible explanatory factor is change in editorial policy. The editorship of *Biology and Philosophy* started with Michael Ruse, who launched the journal in 1986, and then passed on to Kim Sterelny in 2000.[15] Under Ruse, the journal was known to publish historical type papers as well as many thematic issues and brief book notes. The frequency, in the 1980s-1990s, of topics related to evolutionary ethics (#44 Evolutionary ethics, #60 Morality) and evolutionary epistemology (#59 Science) or with a more distinct historical dimension (#2 Nature-Philosophy) fits well with this informal claim. It also fits with Ruse's own recollection of his career

Footnote 13 (continued)

well as most if not all of the topics of categories *G-Socio-normative-issues*, *H-General philosophy of science*, *I-Others* but also many in *E-Network-Entropy-Information* and *F-Cognition-behavior*.

[14] One should always bear in mind that discussion- and special-issues tend to create "bumps" in the diachronic patterns of topics. As can be seen in Fig. 1, discussion issues significantly decreased from the mid 2000s. We have not specifically tracked down special issues (as this is not consistently documented in the metadata we could obtain), but a rough estimate also shows a decrease in frequency over time. This is probably a good thing—at least from our perspective—as special issues tend to artificially create interest in specific topics. Of course, others could argue that special issues foster research and make for more stimulating work. This is another argument that lays outside of our analyses.

[15] Michael Weisberg has been nominated editor since 2017, which is the last year included in the corpus; this change in editorship is therefore unlikely to affect our findings.

as philosopher of biology (Ruse 2016). Sterelny's influence as editor from the 2000s onward likely brought change in this respect, as reflected by the decrease in frequency of these topics in the early 2000s (see Figs. 9, 10, 11). Reflecting on his 17 years as editor of the journal, Sterelny indeed confided his hope that "the center of gravity of the journal changed somewhat as it changed hands" (Sterelny 2017, 2). Whether the changes the topic-modeling revealed are the ones Sterelny seeded or are changes that were bound to happen due to a number of contextual reasons—such as the continued professionalization of the field or general trends in philosophy of science, including its drifting away from historical studies—is difficult to judge, as we could not find any document with explicitly stated editorial policies. Change in editorship nevertheless remains a plausible explanatory factor.

While some of the changes we found agree with other viewpoints on the more general history of the philosophy of biology, others show departure points that could warrant further investigations. Hull, for instance, in 2008, considered the species problem to "continue[…] to tax philosophers and biologists alike" (Hull 2008, 28), yet failed to appreciate that interest in the topic was already declining since the mid 1990s (see Fig. 7). He also estimated that "systematics, reduction, sociobiology, and of all things, fitness [had] ceased being so prevalent", their place being "taken by evolutionary ethics and evolutionary epistemology" (2008, 28–29). On these topics, our results show more nuanced views over the same time-period (see Figs. 7, 9): they agree with a decrease of interest in systematics and reduction (topics #27 Character-taxon, #17 Species, #19 Reductionism) and also in sociobiology (which is part of topic #44 Evolutionary ethics), yet they disagree about evolutionary ethics and evolutionary epistemology that have shown a decrease of presence in the journal rather than an increase (topics #44 Evolutionary ethics, #60 Morality and #59 Science which includes evolutionary epistemology). Hull was also convinced that "ecology, genes, and development" had become increasingly important. Our results concur on genetics and development-related topics (see Fig. 8, topics #64 Mendelian genetics, #12 Genes-DNA, #43 Genes-traits as well as topics #47 Development-Evolution, #68 Development-innateness and #8 Development-morphology), yet the evolution of ecology-related topics appears to be better characterized as fluctuating, with an increase in the 2000s followed by a decrease and an increase again in the mid 2010s (that is to say, after Hull's paper which was published in 2008; see Fig. 7, topics #5 Biodiversity and #10 Ecology).

Conclusion

Revisiting three decades of *Biology and Philosophy* through the lenses of topic-modeling approaches provides a wealth of detailed perspectives, be they about the plurality of topics that philosophers of biology have engaged with, or about the evolution of philosophers' research interests over time. Because these approaches are data-driven and unsupervised, they offer a quantitative and bottom-up view that complements existing historical work on the discipline. They also provide the type of data that contributes an empirical basis for what might otherwise be informal claims about the very topical content of *Biology and Philosophy* or how the journal might have changed over time. A clear advantage of these computational methods is their capability to analyze large corpora of full-text documents, a task that would be extremely tedious if it were to be done by hand. Of course, results must always be interpreted with a clear understanding of the limitations that are inherent to such computational approaches (as with any modeling in general). Yet, bearing in mind these limitations, our analyses have corroborated known episodes of the philosophy of biology and of its evolution, while also shedding light on novel details in terms of topical composition and evolution. While it remains true that evolution by natural selection has been a much prevalent and relatively stable topic in *Biology and Philosophy*, our results document a much broader diversity of topics than is often acknowledged. They also provide detailed insights on their diachronic patterns over the last 30 years—the type of insights that may prompt discussions about the direction the field might (or could) take. After all, philosophy of biology is not just about evolution, in case anyone ever doubted that.

Acknowledgements The authors are grateful to Springer for kindly providing access to *Biology and Philosophy*. The authors also thank Michael Ruse, Phillip Honenberger and Charles Pence for valuable insights on the history of the philosophy of biology as well as Jean-François Chartier for text-mining expertise. Funding from Canada Foundation for Innovation (Grant 34555), Canada Social Sciences and Humanities Research Council (Grant 435-2014-0943) and Canada Research Chair (CRC-950-230795) is gratefully acknowledged. FL acknowledges support from the Fonds de recherche du Québec - Société et culture (Grant 276470)

Authors contribution CM conceived the study, analyzed the results and wrote the manuscript. DP collected/pre-treated the corpus, ran the LDA analyses and contributed to the methodology section. FL contributed to the corpus pre-treatment.

References

Aggarwal CC, Zhai CX (2012) Mining text data. Springer, Berlin

Ayala FJ, Arp R (eds) (2010) Contemporary debates in philosophy of biology. Wiley, Chichester

Ayala FJ, Dobzhansky T (eds) (1974) Studies in the philosophy of biology: reduction and related problems. Macmillan, London

Barron ATJ, Huang J, Spang RL, DeDeo S (2018) Individuals, institutions, and innovation in the debates of the French revolution. Proc Natl Acad Sci 115(18):4607–4612. https://doi.org/10.1073/pnas.1717729115

Bastian M, Heymann S, Jacomy M (2009) Gephi : an open source software for exploring and manipulating networks. In: International AAAI conference on weblogs and social media

Blei DM, Lafferty JD (2009) Topic models. In: Srivastava AN, Sahami M (eds) Text mining: classification, clustering, and applications. CRC, Boca Raton, pp 71–94

Blei DM, Ng AY, Jordan MI (2003) Latent Dirichlet allocation. J Mach Learn Res 3(March):993–1022

Buckner C, Niepert M, Allen C (2011) From encyclopedia to ontology: toward dynamic representation of the discipline of philosophy. Synthese 182(2):205–233. https://doi.org/10.1007/s11229-009-9659-9

Byron JM (2007) Whence philosophy of biology? Br J Philos Sci 58(3):409–422. https://doi.org/10.1093/bjps/axm021

Callebaut W (1993) Taking the naturalistic turn or how real philosophy of science is done: conversations with William Bechtel…[et al.]. University of Chicago Press, Chicago

Chartier J-F, Meunier J-G (2011) Text mining methods for social representation analysis in large corpora. Pap Soc Represent 20(37):1–47

DiMaggio P, Nag M, Blei D (2013) Exploiting affinities between topic modeling and the sociological perspective on culture: application to newspaper coverage of U.S. Government Arts Funding. Poetics 41(6):570–606. https://doi.org/10.1016/j.poetic.2013.08.004
Evans JA, Aceves P (2016) Machine translation: mining text for social theory. Annu Rev Sociol 42:21–50. https://doi.org/10.1146/annurev-soc-081715-074206
Firth JR (1957) A synopsis of linguistic theory 1930–1955. In: Firth JR (ed) Studies in linguistic analysis. Blackwell, Oxford, pp 1–32
Garvey B (2007) Philosophy of biology. Acumen, Stocksfield
Gayon J (2009) Philosophy of biology: an historico-critical characterization. In: Brenner A, Gayon J (eds) French studies in the philosophy of science: contemporary research in France. Springer, Dordrecht, pp 201–212
Godfrey-Smith P (2014) Philosophy of biology. Princeton University Press, Princeton
Grene M, Depew D (2004) The philosophy of biology: an episodic history. Cambridge University Press, Cambridge. https://doi.org/10.1017/CBO9780511819018
Grene M, Mendelsohn E (eds) (1976) Topics in the philosophy of biology. Reidel, Dordrecht
Griffiths P (2008) Philosophy of biology. In: Zalta EN (ed) The Stanford encyclopedia of philosophy, Spring 2018 edn. The Metaphysics Research Lab, Stanford
Griffiths TL, Steyvers M (2004) Finding scientific topics. Proc Natl Acad Sci 101(suppl 1):5228–5235. https://doi.org/10.1073/pnas.0307752101
Griffiths TL, Steyvers M, Tenenbaum JB (2007) Topics in semantic representation. Psychol Rev 114(2):211. https://doi.org/10.1037/0033-295X.114.2.211
Honenberger P (2018) Darwin among the philosophers: Hull and Ruse on Darwin, Herschel, and Whewell. HOPOS 8(2):278–309. https://doi.org/10.1086/698894
Hu Y, Boyd-Graber J, Satinoff B, Smith A (2014) Interactive topic modeling. Mach Learn 95(3):423–469. https://doi.org/10.1007/s10994-013-5413-0
Hull DL (1969) What philosophy of biology is not. Synthese 20(2):157–184
Hull DL (1974) Philosophy of biological science. Pearson College Div, Englewood Cliffs, NJ
Hull DL (2008) The history of the philosophy of biology. In: Ruse M (ed) The Oxford handbook of philosophy of biology. Oxford University Press, New York, pp 11–33
Hull DL, Ruse M (eds) (1998) The philosophy of biology. Oxford University Press, Oxford
Hull DL, Ruse M (eds) (2007) The Cambridge companion to the philosophy of biology. Cambridge University Press, New York
Kampourakis K (ed) (2013) The philosophy of biology. A companion for educators. Springer, Dordrecht
Malaterre C, Chartier J-F, Pulizzotto D (2019) What is this thing called philosophy of science? A computational topic-modeling perspective 1934–2015. HOPOS 9(2):215–249. https://doi.org/10.1086/704372
Matthen M, Stephens C (2007) Philosophy of biology. Elsevier, Amsterdam
Mauri M, Elli T, Caviglia G, Uboldi G, Azzi M (2017) RAWGraphs: a visualisation platform to create open outputs. In: Proceedings of the 12th biannual conference on Italian SIGCHI chapter—CHItaly'17. ACM Press, Cagliari, pp 1–5. https://doi.org/10.1145/3125571.3125585
Mimno D (2012) Computational historiography: data mining in a century of classics journals. J Comput Cult Herit (JOCCH) 5(1):3. https://doi.org/10.1145/2160165.2160168
Murdock J, Allen C, DeDeo S (2017) Exploration and exploitation of Victorian science in Darwin's reading notebooks. Cognition 159(February):117–126. https://doi.org/10.1016/j.cognition.2016.11.012
Nicholson DJ, Gawne R (2015) Neither logical empiricism nor vitalism, but organicism: what the philosophy of biology was. Hist Philos Life Sci 37(4):345–381. https://doi.org/10.1007/s40656-015-0085-7
Peirson BRE, Bottino E, Damerow JL, Laubichler MD (2017) Quantitative perspectives on fifty years of the journal of the history of biology. J Hist Biol 50(4):695–751. https://doi.org/10.1007/s10739-017-9499-2
Peldszus A, Stede M (2013) From argument diagrams to argumentation mining in texts: a survey. Int J Cognit Inform Nat Intell (IJCINI) 7(1):1–31. https://doi.org/10.4018/jcini.2013010101
Pence CH, Ramsey G (2018) How to do digital philosophy of science. Philos Sci 85(5):930–941. https://doi.org/10.1086/699697
Pradeu T (2017) Thirty years of biology and philosophy: philosophy of which biology? Biol Philos. 32:149–167. https://doi.org/10.1007/s10539-016-9558-7
Pritchard JK, Stephens M, Donnelly P (2000) Inference of population structure using multilocus genotype data. Genetics 155(2):945–959

Ramsey G, Pence CH (2016) EvoText: a new tool for analyzing the biological sciences. Stud Hist Philos Biol Biomed Sci 57(June):83–87. https://doi.org/10.1016/j.shpsc.2016.04.003
Rosenberg A, Arp R (eds) (2009) Philosophy of biology: an anthology. Wiley, Malden, MA
Rosenberg A, McShea DW (2008) Philosophy of biology: a contemporary introduction, 1st edn. Routledge, New York. https://doi.org/10.4324/9780203926994
Ruse M (1973) The philosophy of biology. Hutchinson, London
Ruse M (1988) Philosophy of biology today. SUNY Press, Albany
Ruse M (1989) David Hull through two decades. In: Ruse M (ed) What the philosophy of biology is: essays for David Hull. Kluwer Academic Publishers, Dordrecht, pp 1–15
Ruse M (ed) (2008) The Oxford handbook of philosophy of biology, 1st edn. Oxford University Press, Oxford
Ruse M (2016) L'état, c'est Moi : fifty years of history and philosophy of evolutionary Biology. Theor Biol Forum 1–2(109):111–122. https://doi.org/10.19272/201611402008
Sarkar S, Plutynski A (eds) (2008) A companion to the philosophy of biology. Blackwell Pub, Malden, MA
Sober E (1999) Philosophy of biology. Westview Press, Boulder, CO
Sober E (ed) (2006) Conceptual issues in evolutionary biology, 3rd edn. MIT Press, Cambridge, MA
Srivastava AN, Sahami M (2009) Text mining: classification, clustering, and applications. CRC Press, Boca Raton
Sterelny K (2017) Farewell. Biol Philos 32(1):1–3. https://doi.org/10.1007/s10539-016-9559-6
Sterelny K, Griffiths P (1999) Sex and death: an introduction to philosophy of biology. University of Chicago Press, Chicago, IL
Swanson R, Ecker B, Walker M (2015) Argument mining: extracting arguments from online dialogue. In: Proceedings of the 16th annual meeting of the special interest group on discourse and dialogue, pp 217–226. https://doi.org/10.18653/v1/W15-4631
Turney PD, Pantel P (2010) From frequency to meaning: vector space models of semantics. J Artif Intell Res 37:141–188. https://doi.org/10.1613/jair.2934
Widdows D (2004) Geometry and meaning. CSLI Publications, Stanford

Affiliations

Christophe Malaterre[1,2] · Davide Pulizzotto[3] · Francis Lareau[4]

[1] Département de philosophie, Université du Québec à Montréal (UQAM), 455 Boulevard René-Lévesque Est, Montréal, QC H3C 3P8, Canada

[2] Centre interuniversitaire de recherche sur la science et la technologie (CIRST), Université du Québec à Montréal (UQAM), 1205 rue Saint-Denis, Montréal, QC H3C 3P8, Canada

[3] Département d'études littéraires, Université du Québec à Montréal (UQAM), 405 rue Sainte-Catherine Est, Montréal, QC H3C 3P8, Canada

[4] Département d'informatique, Université du Québec à Montréal (UQAM), 201, av. Président-Kennedy, Montréal, QC H3C 3P8, Canada