Author Pre-press document
Final version available at http://dx.doi.org/10.1007/s11367-016-1153-2.

Implicit Prioritization in Life Cycle Assessment: Text Mining and Detecting Metapatterns in the Literature

Emily Grubert
Emmett Interdisciplinary Program in Environment and Resources
Stanford University
473 Via Ortega Suite 226, Stanford CA, 94305 USA
Email: gruberte@stanford.edu; phone: 510-545-3796

Abstract

Purpose Life cycle assessment aims to evaluate multiple kinds of environmental impact associated with a product or process across its life cycle. Objective evaluation is a common goal, though the community recognizes that implicit valuations of diverse impacts resulting from analytical choices and choice of subject matter are present. This research evaluates whether these implicit valuations lead to detectable priority shifts in the published English language academic LCA literature over time.

Methods A near-comprehensive investigation of the LCA literature is undertaken by applying a text mining technique known as topic modeling to over 8,200 environment-related LCA journal article titles and abstracts published between 1995 and 2014.

Results and Discussion Topic modeling using MALLET software and manual validation shows that over time, the LCA literature reflects a dramatic proportional increase in attention to climate change and a corresponding decline in attention to human and ecosystem health impacts, accentuated by rapid growth of the LCA literature. This result indicates an implicit prioritization of climate over other impact categories, a field-scale trend that appears to originate mostly in the broader environmental community rather than the LCA methodological community. Reasons for proportionally increasing publication of climate-related LCA might include the relative robustness of greenhouse gas emissions as an environmental impact indicator, a correlation with funding priorities, researcher interest in supporting active policy debates, or a revealed priority on climate versus other environmental impacts in the scholarly community.

Conclusions and Recommendations As LCA becomes more widespread, recognizing and addressing the fact that analyses are not objective becomes correspondingly more important. Given the emergence of implicit prioritizations in the LCA literature, such as the impact prioritization of climate identified here with the use of computational tools, this work recommends the development and use of techniques that make impact prioritization explicit and enable consistent analysis of result sensitivity to value judgments. Explicit prioritization can improve transparency while enabling more systematic investigation of the effects of value choices on how LCA results are used.

Keywords: life cycle assessment, prioritization, weighting, text mining, topic model, environmental impact, methodology

1 Introduction

Life cycle assessment (LCA) is a tool developed largely to facilitate decision making on the basis of information that is as complete as possible (see e.g. Nash and Stoughton 1994, Hertwich and Hammit 2001). Completeness is defined in two dimensions: 1) examining every stage of an analyte's existence, such that all processes involved in the creation, use, and disposal of the product or process are considered; and 2) examining all impacts, such that the effects of a process on a comprehensive list of environmental, social, and economic indicators are considered, measured, and reported. One important task for LCA researchers is to ensure that decision-makers and other users of LCA data have maximally complete information about all types of impacts associated with the functional unit being considered. Clearly articulating where data gaps and other concerns lead to simplifications or exclusions that might result in the intentional or unintentional prioritization of a particular impact category, indicator, or interpretation is critically important to preserving the multidimensional value of LCA.

The focus of the research presented here is whether the LCA research literature reflects a field-scale trend toward implicit prioritization of certain impact categories, which supports a broader inquiry as to whether implicit prioritization is a concern and how it can be addressed. The computer-aided investigation of over 8,200 English language LCA-related journal article abstracts and titles published between 1995 and May 2014 presented here indicates that environmental focus areas have indeed shifted over time. This shift, which is difficult to detect at the individual article level, highlights the considerable power of the scientific community to mediate the information decision-makers access. This article introduces a text mining method called topic modeling to the LCA community and demonstrates that digital tools are relevant and useful for enabling the investigation of field-scale patterns that influence the application of powerful scientific tools like LCA.

1.1 Implicit and explicit prioritization in LCA

A major motivation for investigating LCA at a field scale is that patterns of decisions made in the application of this still-developing field can create subtle shifts in the way that scientists ask and answer questions of substantial societal importance over time. One of LCA's greatest strengths is its methodical examination of diverse environmental impacts: this work investigates whether that strength is potentially weakened by a field-scale trend toward reduced diversity in the impact categories actually being assessed. Assessing multiple types of potential harm helps users identify priority areas for mitigation, recognize perhaps unexpected impacts, and otherwise think comprehensively about a situation. Consistently investigating all of the impact categories recommended in method-defining documents like the ISO standards helps protect a greater degree of objectivity in the LCA method, with the goal of exhaustive examination of environmental impact without valuation that is highly valued by many researchers (see e.g. Marsmann et al 1999; Owens 1997; Schmidt and Sullivan 2002).

This interest in objectivity is often observed in the form of a proscription of weighting schemes, including the explicit statement in ISO 14044 that

> "weighting, shall not be used in LCA studies intended to be used in comparative assertions intended to be disclosed to the public" (ISO 2006).

Concern about weighting (see e.g. Finnveden et al 2009) is grounded in part in LCA-based decision-making's necessary engagement with heavily values-based issues like whether a given improvement in air quality is worth some level of water quality degradation. Such issues are clearly context-dependent and likely better addressed outside the expert scientific community once appropriate measurement, characterization, and normalization have occurred (Grubert under review). However, given the challenges associated with making such tradeoffs, there are long-standing acknowledgments in the community that providing a robust, science-based, transparent method of comparing dissimilar impacts that is integrated with decision-maker needs is desirable (see e.g. Hertwich et al 1997; Zhang et al 2010(a), (b); Plevin et al 2014). These competing interests in leaving the value judgments to decision-makers while also wanting to provide robust and helpful information to decision-makers who might be unfamiliar with the complex underpinnings and outputs of LCA imply that an explicit, rigorously designed, and transparent weighting scheme is preferable to the difficult to detect, implicit weighting schemes that can emerge as patterns of researcher decisions.

There are many necessary points of valuation and prioritization within the LCA framework independent of whether an explicit weighting scheme is ultimately applied to the results: truly objective LCA is not possible. Hauschild (2005) notes that goal and scope definition is "obviously" a value choice: as scientists, our choices of what to study inherently apply values. Similarly, choices about which impact categories to include are implicit weighting schemes (Souza et al 2014) that essentially weight included categories at 1 and excluded categories at 0. This issue, too, has been recognized by many authors. For example, Ayres (1995) notes a choice to prioritize energy over toxics; Zhang et al (2010(a), (b)) note that LCA tends to exclude some ecosystem services from consideration; and Scanlon et al (2014) note that a decision to focus on environmental impacts outside the workplace might systematically lead to a shifting of impacts from the external world to the work environment, disproportionately harming workers.

Given that value judgments are inherent to scientific analysis, including LCA, this work posits that the direction of implicit prioritization in LCA at a field scale is important to recognize and acknowledge. One major reason that implicit prioritization is important is that value choices made at a methodological level are often self-reinforcing because authors are expected to cite existing literature and often to prioritize more recent work. Particularly since the number of articles published on LCA has increased dramatically over time (Online Resource Figure S1), a field-scale shift in focus is likely to become ingrained in the literature. Since decision-makers often turn to contemporary academic literature as a source of accurate, up-to-date scientific information, focal shifts could affect what types of policies attract the attention of stakeholders (see e.g. Miller and Neff 2013). This analysis particularly investigates whether the LCA literature has undergone a topical shift over time because of the uneven relevance of different impact categories to stakeholders. For example, certain communities might not be motivated by discussion of anthropogenic climate change but might welcome discussion of impacts on water availability as a result of the same project: a focus on climate over water in the LCA literature might thus reduce opportunities for cooperation and discourse. Thus, it is relevant to ask whether as a community, LCA practitioners are systematically prioritizing some impacts over others, what that means, and how we as a community can improve the scientific basis for judgments that need to be made.

1.2 Measuring the shift in LCA analytical focus

As a method, LCA has come of age in roughly the same period as concern over climate change. LCA was formalized under the ISO 14000 series in 1996 (ISO 1996); the highly impactful Rio conference on climate change took place in 1992. Attention to both have increased over time. Recent LCA work in particular seems to emphasize greenhouse gas emissions as an indicator, perhaps at the expense of other environmental indicators (McManus and Taylor 2015): whether this apparent trend is real is the focus of this work. The importance of a potential climate emphasis is most saliently that a value judgment is being made without explicit justification when greenhouse gas emissions are investigated and indicators for other impact categories are not. Other authors have noted that a single-issue focus on greenhouse gases can be problematic from a sustainability perspective. For example, Turconi et al (2013) note that based on their assessment of 167 case studies focused on electricity generation LCA, greenhouse gas emissions are an inappropriate single indicator to assess environmental performance. Guinée et al (2011) note in their review of the past, present, and future of LCA that while some recent studies are increasingly sophisticated both in their modeling and in their inclusion of impact categories, others narrowly focus on carbon footprinting and can miss important other sustainability-related effects. Also, interestingly, Masanet and Chang (2014) note in a study of the sustainability priorities of potential LCA practitioners that climate is not the top priority for the next generation of researchers. Sustainable energy and water both rank higher than climate, which less than half of respondents rate among their top three named priorities.

The contribution of this research is to build on the extensive methodological refinements and critiques of value judgments and prioritization in environmental LCA partially described here by applying a broader lens to questions of how the community is evolving. Specifically, this research uses computer-aided text mining techniques to elicit information about topical focus in LCA on a nearly comprehensive dataset, adding to the intuition of many that there are potentially meaningful value judgments being made at the research community level. That is, this work investigates whether similar value judgments are being made frequently enough to create a bias in the literature.

To answer the question of whether and how topics treated in the academic LCA literature have changed over time, this work uses topic modeling to characterize a corpus of over 8,200 English-language abstracts and article titles related to environmental LCA written between 1995 and 2014. This is considered near-exhaustive based on related work by e.g. Guinée et al (2011) and McManus and Taylor (2015) and on investigation of articles dating to as early as 1900 for relevance.

1.3 Text mining and topic modeling

Text mining is an increasingly popular term for a set of computational tools applied to words and phrases in documents as the units of analysis. Text mining tools like word frequency analysis, text clustering, sentiment analysis, and topic modeling find applications in marketing (Sullivan 2001), security (Corney et al 2002; Gegick et al 2010), policy (Talamini et al 2012), and many academic fields, including life cycle assessment and other environmental fields (Kostoff et al 2008; Neff and Corley 2009; Altaweel and Bone 2012; Zamagni et al 2012; Varsara et al 2013). While many of these types of analysis have been carried out manually in the past, the advent of computer aids has made much larger studies possible and far more accessible.

One application of text mining, and the one primarily used in this work, is known as topic modeling. Topic modeling is an increasingly common tool in a wide range of fields: Blei, Ng, and Jordan's (2003) original paper introducing the tool has been cited over 10,000 times. As a technique that focuses on making information available from very large bodies of text, topic modeling is well suited to literature reviews and metaanalysis of specific fields where a goal is to understand themes in a collection of documents (known as a corpus) that is too large to evaluate manually. Benefits of computationally-aided analyses include greater reproducibility, transparency, and scale relative to traditional review techniques (Grubert and Siders under review). For this study, for example, a manual review of how research attention to various impact categories in English-language LCA has changed over the last twenty years would have required the thorough reading and consistent coding of over 8,200 documents, which is a relatively small corpus (see e.g. Neff and Corley 2009 for an example of a 160,000 document investigation in ecology).

This study uses unsupervised topic modeling, which means that the algorithm is not primed with content awareness (for example, in the form of a specific word list to seek in the corpus). The use of unsupervised techniques is important for this investigation into the state of the LCA literature, as it means that the discovery process is not pre-informed by a hypothesis. Supervised text mining techniques, where the analyst provides the algorithm with examples and focus areas, can be extremely valuable when researching a specific topic in context, but unsupervised techniques like the topic models used in this work are more useful when the tool is being used to investigate a hypothesis. By allowing a naive algorithm to discover themes across a corpus without direction, the research presents an independent measurement of the attention paid to each topic by authors of documents in the corpus that can then be interpreted by a human with content knowledge sufficient to derive insights. While this interpretation step is subjective, the basis for interpretation (the model output) is far more open to examination and critique than with a fully manual review. More detail on topic modeling is given in Methods (2.2.1).

2 Methods

2.1 Data collection

The dataset used in this work is a collection of over 8,200 scientific journal article titles and abstracts, classified by year of publication, published between January 1995 and May 2014. The titles and abstracts were downloaded from the Web of Knowledge interface using the key word "topic = 'life cycle,'" which captures variants like "life cycle cost," "life cycle inventory," "life cycle analysis," and "life cycle assessment." (Note that "topic" in Web of Knowledge parlance is not the same "topic" found in a topic model.)

As the goal of this work was to examine topics in environmental life cycle assessment research, filters were applied to restrict results to English-language publications in fields likely related to environment and energy (Online Resource Table S1). While the filters are permissive and likely overinclusive of nonenvironmental social science research, they were used to exclude papers investigating biological life cycles, as the phrase denotes significantly different work in the biological and health sciences versus environmental science (usually pertaining to the life cycle of an organism). Results from included fields were also manually inspected for topicality. "Ecology" and "psychiatry" were excluded in this second phase due to a preponderance of research on biological life cycles and drug life cycles, respectively. The full dataset, classified by year of publication, is available online (Grubert 2014).

Web of Knowledge was chosen as the data source for this text mining work due to its inclusion of major environment-focused journals, despite the lack of access to bulk full-text downloads. Abstracts and titles are considered representative of article topics due to their purpose and summary-oriented format. While both JSTOR and Elsevier enable bulk full text downloads (at least in bag-of-words form), both also lack key environment journals.

2.2 Topic modeling

2.2.1 How topic modeling works

Topic modeling works by assigning each individual word in a document to a list of words and the probability that they occur in a given cluster called a topic, thereby creating a many-to-many mapping between a corpus of documents and a set of topics, where "k" denotes the number of topics chosen by the user. The value of k is ultimately a subjective choice, though the choice can be guided by sensitivity analyses (e.g. performing topic modeling for a variety of choices of k and selecting the value or values that provide the most insight based on semantic interpretation; this study investigated k = [2,10]). Additionally, the choice of k is influenced by the goal of a study: for a highly dispersed corpus with limited self-similarity where the goal is to find unusual topics, k might be 100 or higher; for a tighter corpus where the goal is to identify broad themes, as with this study, k might be less than 10. Each of the k topics is a list of all the words in a corpus in descending order of probability that a given topic will generate each word, based on a statistical model described below. Topics are usually hand-classified, or labeled, based on the highest probability words and the purpose of the analysis. For example, a topic with "rose, iris, poppy, mallow, daisy" as its five highest probability words might be classified as "flower-related," "nouns," or "living organisms" depending on the goal of the analysis. Since the assignment of labels to topics is both important and subjective, it is good practice to make the top n words for each topic available for inspection.

Other user choices include choice of dataset and choice of topic modeling tool. This study uses journal article abstracts for studies using a specific method in a specific time period because the goal is to determine whether analytical priorities in the academic community have changed as LCA has become a more popular technique. The tool used for this research is a common technique called Latent Dirichlet Allocation, or LDA, using Gibbs sampling. While other techniques have proven to have significant advantages in certain situations, basic LDA works well in this application in part due to the similarity of the goal here and that for which LDA was initially developed. One of the original LDA test corpora was a set of 5,225 journal abstracts related to C. elegans (Blei et al 2003), and later work focused on mapping articles from *Science* in JSTOR (Blei and Lafferty 2006, Blei and Lafferty 2007).

LDA is a mixture model where individual words are assumed to be part of a document because of the topics the document addresses. That is, given that a document is characterized by some number of topics k, certain words comprising each topic are expected to occur with some probability. Thus, by examining which words actually do occur in a document, inferences about which topics are present in that document can be made. Since topics are not user defined, the method must also generate topics given a list of words from a corpus. Very common words that appear frequently in almost every topic (like "and," "for," "the," etc.), often referred to as "stop words," are used to answer questions of authorship, style, and subconscious gendered language but are typically removed when the goal is to uncover content topics (see e.g. Jockers and Witten 2010), as is the case here.

Topic generation in LDA is typically done via estimation methods like Gibbs sampling, a common technique also applied in this work. Gibbs sampling proceeds by fitting a model that assumes documents can be described by the bag-of-words model, which discards word order but preserves the original association of words and their frequencies to a particular document. Then, words are evenly divided across k topics at random, where the value of k is an analytical choice as described above. Reallocation of individual words from these k random topics into more coherent topics is based on two probabilities: the probability that a given word appears in a topic as it currently exists (that is, for each k_i, what is p(word) in k_i?) and the probability that a word from a given topic as currently assigned appears in the document of interest (that is, for each k_i and document x_j, how much of x_j belongs to k_i?). Based on these probabilities, a word is either moved to a different topic or retained in its currently assigned topic. When this sampling process has been repeated sufficiently often that randomly selected words are no longer being reallocated, the topics are considered stable. For an intuitive description of the method, see Jockers 2011.

2.2.2 Applying Topic Modeling to the Academic LCA Literature

The dataset of over 8,200 English language journal abstracts obtained by the method outlined in Section 2.1 was topic modeled using LDA in MALLET (MAchine Learning for LanguagE Toolkit) (McCallum 2002) via the Topic Modeling Tool UI (Google 2011). MALLET was applied in two contexts: first, to the entire collection of documents published between 1995 and 2014 and classified by publication year; then, to subcorpora comprised of abstracts from specific topical or general interest journals that published large numbers of LCA articles between 1995 and 2014. The first investigation addressed how focus changed through time in nearly the entire LCA literature, while the second addressed whether a changing focus was field-wide versus potentially driven by publications in more topically focused journals.

For all topic models presented in this work, MALLET was run on article titles and abstracts grouped by year using 1,000 iterations, a topic threshold of 0.05, and an augmented stopwords list defined as a common list (Salton 1971) plus context-specific additions based on common phrases found in journal abstracts and appearing across multiple topics without adding significant understanding: ‘Elsevier,’ ‘life,’ ‘cycle,’ ‘assessment,’ ‘lca,’ ‘rights,’ ‘reserved,’ ‘paper,’ ‘study,’ and ‘results.’ MALLET returned the top 20 words for each topic. Stemmers were not used here due to a desire to preserve subtle differences in meaning between words like “transport” and “transportation” (where transport often refers to the movement of contaminants, while transportation often refers to systems for moving people and products through society), “gas” (as in greenhouse gas or natural gas) and “gasoline” (as in oil-derived fuel), and others. The relatively consistent use of language in English-language journal abstracts and manual inspection support the idea that stemmers do not provide major other benefits in this context.

Hand-labeling of the topics proceeded by first investigating top words by topic for MALLET output across multiple runs for each k. This repetition was performed to ensure robustness and to increase familiarity with the types of words entering each topic. Individual documents showing high proportions for given topics were also manually inspected to increase intuition. Prior to labeling, over 300 environment and energy-related life cycle assessment articles were read, which also contributed to intuition about how words correspond to research topics during hand-labeling. Entirely objective hand-labeling (and particularly naming) of topics is not possible, and top words for each topic are included in the Online Resource, Tables S2 to S8, alongside the labels used for this research for transparency (three-topic model words are given in Table 1 as an example).

Full 1995-2014 corpus, classified by year For the full corpus, abstracts ($n = 8{,}239$) were sorted and divided by year of publication such that MALLET associated each abstract to an umbrella “document” containing all abstracts published in a given year ($n = 20$). MALLET was run in detail for $k = [3,7]$. Additional low resolution trials were performed for $k = 2$, 8, 9, and 10 and deemed nonadditive to the $k = [3,7]$ results. In particular, $k = 2$ did not produce meaningful topic separation, and $k > 7$ produced overdefined topics that hand-labeling interpreted as too similar to differentiate given the relatively low resolution of this investigation into broad categories of environmental impact. While the results of a single, specific run are presented in this work, MALLET was run multiple times for each k and topics were hand-labeled each time and compared to ensure robustness of the result based on seed value. Top topic words and labels for the overall corpus can be found in Table 1 and in the Online Resource, Tables S2 to S5.

Subcorpora, topical and general interest journals Subcorpora comprised of abstracts published in the three journals with the largest numbers of published LCA articles between 1995 and 2014 were topic modeled separately to test the resolution of topical shift in the LCA literature. *The International Journal of Life Cycle Assessment* (*JLCA*, n = 886) is an LCA method-specific journal, while the *Journal of Cleaner Production* (*JCP*, n = 694) and *Environmental Science and Technology* (*ES&T*, n = 456) are general interest environment journals. For each subcorpus, a five-topic model was run using MALLET and the settings described above: the number of articles published by each journal each year is presented in Table 3, and top topic words and labels for the subcorpora can be found in the Online Resource, Tables S6 to S8. The choice of k = 5 for presentation is based on results of low resolution trials for values of k between 3 and 10. To reduce the risk of confirmation bias, topics were charted and analyzed for trends before word probabilities were viewed and labels were assigned.

3 Results

3.1 Full 1995-2014 corpus, classified by year Topic proportions and highest probability words for LCA abstracts and titles by publication year are presented in Figures 1-5, Table 1, and the Online Resource, Tables S2 to S5, also available interactively online at lcatopics.emilygrubert.org. Topic tags are chosen subjectively based on these highest probability words and used for labeling and analysis. Publishing the highest probability words is intended to add transparency to this process that likely exceeds transparency associated with a traditional literature review, where conclusions are similarly drawn based on author judgment and experience but are harder to replicate.

A clear tradeoff between the topics labeled “health” and “climate change” is visible, most evidently for the simple k = 3 model (Figure 1, Table 1). This trend persists for k = [4,7], with additional topics primarily adding to resolution on the topic labeled “systems analysis” for the k = 3 model (Figures 2-5). As k increases, topics related to waste management (including pollution management in general), risk assessment (linked to both health risks and overall ecological risks that might have climate-related drivers), and more specific energy and environmental systems emerge. Additionally, at k = 7, a clear distinction between a human and nonhuman health topic emerges.

3.2 Subcorpora, topical and general interest journals Most common words for topics in k = 5 models of single-journal subcorpora and their proportional contributions to the corpora are presented in the Online Resource (Tables S6 to S8 and Figures S2 to S4).

The most LCA-associated journal in the corpus, *JLCA* (n = 886), reflects relatively well balanced topics in its titles and abstracts (Figure S2). The topic with the greatest growth, labeled "balanced impacts," includes the words "land," "social," "water," "carbon," and "recycling" in its top 20, indicating that *JLCA* articles have tended not to emphasize a single impact category on average. There is no topic that appears heavily focused on climate change, but both topics "balanced impacts" and "energy" include clearly climate-associated words (e.g. "carbon," "greenhouse"). The topic with the greatest decline, labeled "waste and health," is largely comprised of formulaic LCA words (e.g. "method" and "scope"), with some health-associated words ("case," "characterisation," "health," "damage") as well. This topic is only weakly health associated, but it is notable that health words (other than "ecotoxicity" in the "water" topic) do not appear in the top 20 for any of the other four topics. Given the journal launch date of 2001 and trends in the overall corpus, this result is consistent with findings from the overall corpus.

The two generalist journals investigated for this journal-specific analysis, *JCP* (n = 694) and *ES&T* (n = 456) also exhibit less coherent topics than the overall dataset, largely due to small numbers of articles per year (Table S2). Declining topics in *JCP* are labeled "waste" and "methods," while ascending topics are labeled "balanced impacts" and "climate" (Figure S3). No specifically health-associated words were identified in the top 20 words of the five topic model presented here. Indeed, the most human- or social-focused word in the topics is "economic," found in both the "waste" and "balanced impacts" category. Thus, the growth of climate as a topic is observed in this generalist journal, and health is not present. Given the first publication date of 2002 and the trends displayed in the overall corpus, this lack of focus on health impacts is consistent with other results. *ES&T* also exhibits trends consistent with the overall corpus results, with a longer record of LCA (the first LCA articles in *ES&T* were published in 1993). Specifically, declining topics are labeled "health and economy" and "risk modeling," while ascending topics are labeled "water" and "climate" (Figure S4). Unlike *JLCA*, *JCP* and *ES&T*

both returned topics in a k = 5 model that are fairly coherently climate-oriented. This result was tested by re-running the topic model several times and was consistently encountered.

4 Discussion

4.1 The rise of greenhouse gas emissions as an environmental impact indicator

Overall, topic modeling supports the hypothesis that users of LCA have shifted focus from health to climate change over the past 20 years, and the proportional emphasis on climate appears to be climbing. Modeling for multiple values of k shows that the result is robust even with more nuanced topics, also emphasizing the value of different model resolutions for clusters whose quality is based on semantic interpretation. Notably, the health to climate shift is also weakly observable at the journal level for high volume generalist journals in particular, indicating that the trend is not exclusively attributable to high publication rates in climate-oriented journals. Of the 8,239 articles analyzed for this work, only 73 (0.9%) were published in journals with “clima*,” “greenhouse,” or “carbon” in the title. Similarly, none of the top 10 highest volume LCA-publishing journals (accounting for 42% of the dataset) include a specific focus on climate in their aims and scope, though two do mention climate alongside other impacts (*Journal of Cleaner Production*, as “renewable energies and other low-carbon technologies and products” and *Applied Energy*, as “social and economic impacts of energy policies and usage, including climate change mitigation and other environmental pollution reduction”). This result suggests that English-language academic LCA publications reflect an implicit value judgment that privileges climate change impacts above other types of environmental impact.

This increased emphasis on climate in academic LCA publications could derive from several sources, ranging from analytical convenience to the values held by the scientific community. One major question is whether the change is driven by LCA-focused researchers prioritizing climate versus climate-focused researchers adopting LCA as a useful method. While many of the same conclusions might follow in either case, including a recommendation for more explicit frameworks for supporting LCA-based decision-making, the implementation approach might be quite different. For example, if the LCA community is already using more standardized approaches, the goal might be broader dissemination of training about LCA as a method; if the LCA community is itself focusing on climate specifically, the goal might be a more detailed investigation of why climate has taken precedence.

In fact, analysis of topical trends in *JLCA* versus *JCP* and *ES&T* indicates that the trend towards climate precedence is likely driven by articles in journals not focused specifically on LCA as a method. *JLCA* is the only one of the subcorpora analyzed that does not seem to include the rise of a recognizable and coherent “climate” topic over time, which implies that climate change is more frequently treated as an independent impact category outside the most specialized LCA journal. While lower volume journals were not modeled in depth due to their low numbers of articles per year, low resolution topic modeling shows that nonclimate specialty journals (e.g. *Environmental Toxicology and Chemistry*) also tend not to display the increasing climate focus over time.

Given that LCA’s primary journal does not display a strong shift toward climate and none of the high-volume LCA-publishing journals have a specific climate focus, the source of the trend encountered in this work is somewhat unclear. Reasons for this focus could include the relative robustness of greenhouse gas emissions as an environmental impact indicator, a correlation with funding priorities by major grantors, researcher interest in supporting active policy debates, or a revealed priority on climate versus other environmental impacts in the scholarly community.

Climate change potential has certain analytical characteristics as an impact indicator that increase confidence in results. Specifically, unlike most other midpoint indicators, greenhouse gas emissions are generally considered, at least to first approximation, to have the same effect on the environment regardless of when and where they are emitted. While the effect of a factory’s release of a heavy metal depends on many variables, including proximity to populations, proximity to water, specific weather patterns during the period of emission, etc., the emission of greenhouse gases has a fairly well understood, context-independent effect on the environment. Further, the magnitude of such emissions can often be robustly estimated from relatively high quality data based on well understood physical principles like combustion. Also, while the exact endpoint implications of climate change are not well known, it is generally agreed that each midpoint unit of greenhouse gas has roughly the same endpoint effect on the global environment. Since results about the total greenhouse gas footprint of a product therefore tend to be relatively robust, and since article titles and abstracts tend to focus on a study’s strongest conclusions, this title and abstract-based analysis might also be overcapturing climate results.

Exogenous reasons for the surge in climate-related work include demand for climate-focused work, particularly in light of major ongoing international efforts to define climate policy. This demand is potentially reflected as funding availability, which in turn is part of a feedback loop where the quality and quantity of related work both creates and reflects prioritization by funding agencies. For example, McManus and Taylor (2015) argue that the high-volume debate over the climate impacts of bioenergy and biofuel, largely in support of policy change, is visible in the overall LCA literature.

Perhaps most intriguing is the possibility that the focus on climate actually reflects priorities in the scientific community. Scientists might themselves be motivated to work on climate issues specifically, with the publication record reflecting the aggregated priorities of many scientists. For example, Neff describes the way that ecologists choose research topics as potentially motivated by interests in doing research that compels change associated with a problem, informs active policy debates, builds theory, and informs professional action (Neff 2011). Thus, it is quite possible that a community focus on climate reflects scientists' personal feelings that climate change is the most important environmental concern at present, which would suggest that the rise of climate in LCA reflects the priorities and values of the scientific community. This explanation is especially interesting because it suggests that a technique like topic modeling could be used to "survey" the scientific community about social priorities, contributing to a fuller understanding of the "situated knowledge" that Freidberg (2015) and others have described.

4.2 Potential sources of error

Several conditions unrelated to prioritization of climate and deprioritization of health could be contributing to the rise of climate and fall of health as dominant topics in LCA. First, the earlier years have significantly fewer abstracts than the later years, which means that the signal could be skewed by a small sample size. However, the continuation of the trend year-by-year and the fact that this sample is essentially the entire population of LCA-related abstracts for a given year suggests that this is not a major problem. Similarly, the use of a permissive filter in selecting life cycle articles related to energy and environment rather than other topics could be skewing the results, particularly when $n < 100$ and the inclusion of only a few unrelated articles represents a driver of a few percent. Given that the sample is overly inclusive of nonenvironmental social science research, the signal of increased focus on climate is robust, as the filter permissiveness would be expected to dampen the climate signal. However, the signal of decreased focus on health might in reality be less strong than observed, as some early unrelated articles on e.g. family life cycles might be expected to discuss health topics. Even given this potential inflation of health topics in the earlier years, when LCA was less common, the "health" topic is quite large in those years, suggesting that most of the signal comes from relevant articles.

A trickier potential issue is whether the topic model is actually capturing a latent prioritization or merely the most significant result of the paper as revealed in the title and abstract. It is conceivable that a full-text analysis of papers included in the corpus might reveal much broader scopes, but only the climate-related impacts were considered significant enough to report in the abstracts due to e.g. policy attention or earlier resolution of other types of environmental impacts. However, manual reading of many LCA article full texts, in addition to earlier comments by e.g. Turconi et al (2013) and Guinée et al (2011), suggest that this is not the case. Even if the abstracts are reporting significantly climate-skewed versions of broader findings in LCA articles, the interpretation of climate results as the most important is a relevant piece of information about how LCA results are being presented.

5 Conclusions

5.1 Society and Valuation in LCA

LCA is often viewed explicitly as a decision-making tool for environmental policy. In order to make a decision, decision-makers need both information about different types of impacts and a way to make tradeoffs among dissimilar impacts—a preference function. The value of LCA will be greatly enhanced if the framework for making tradeoffs is as strong and clear as the framework for assessing impacts, and enhancing the framework for tradeoffs first requires acknowledging and considering the many implicit valuations that are built into LCA as it currently exists. This work shows that over the last 20 years, the focus of LCA publications has identifiably shifted. The main locus of this shift appears to be in publications focused on using rather than developing LCA as a tool. Thus, an effort from the LCA methodological development community to ensure that users have easy access to frameworks that help to make value judgments more visible could enhance LCA's value as a tool.

5.2 Recommendations and Future Work

The LCA community seems ready to develop and use explicit prioritization schemes that can help elucidate which value judgments are being made and can support decision-making that is as strong as the LCA work, as opposed to defaulting to an intuitive choice (Johnsen and Lokke 2013). At minimum, LCA authors should describe their analytical and methodological choices, whether by explaining why they make certain choices or by explaining what measures they put in place to minimize implicit weighting.

In addition, LCA could benefit from a standardized set of prioritization archetypes to apply in sensitivity analyses and as a way to add nuance to tradeoff discussions. Archetypes based on real, community-sourced priority patterns can add social scientific rigor to a very challenging problem in LCA-based decision-making by providing evidence-based examples of how a value system might affect a decision. Ongoing work uses text mining techniques applied to temporally and spatially specific community-based corpora like news articles, meeting minutes, interview transcripts, and other sources, with validation from traditional survey work, to measure priority patterns and generate such archetypes (Grubert under review).

6 Acknowledgments

Thanks to Mark Algee-Hewitt, Justin Grimmer, Jason Heppler, Stanford's Environmental Assessment and Optimization group, and E-IPER students for coaching, guidance, and feedback on this project. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-114747. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation. The final publication is available at Springer via http://dx.doi.org/10.1007/s11367-016-1153-2.

7 References

[1] Altaweel M, Bone C (2012) Applying Content Analysis for Investigating the Reporting of Water Issues. Comput Environ Urban Syst 36.6:599–613
[2] Ayres RU (1995) Life Cycle Analysis: A Critique. Resour Conservation Recycl 14.3–4:199–223
[3] Blei D, Lafferty J (2006) Correlated Topic Models. Adv Neural Inf Process Syst 18:147
[4] Blei DM, Lafferty JD (2007) A Correlated Topic Model of Science. Ann Appl Stat 1.1:17–35
[5] Blei DM, Ng AY, Jordan MI (2003) Latent Dirichlet Allocation. J Mach Learn Res 3:993–1022

[6] Corney M, de Vel O, Anderson A, Mohay G (2002) Gender-Preferential Text Mining of E-Mail Discourse. Proc 18th Annu Computer Secur Appl Conf doi:10.1109/CSAC.2002.1176299
[7] Finnveden G, Hauschild MZ, Ekvall T, Guinée J, Heijungs R, Hellweg S, Koehler A, Pennington D, Suh S (2009) Recent Developments in Life Cycle Assessment. J Environ Manag 91.1:1–21
[8] Freidberg S (2015) From behind the Curtain: Talking about Values in LCA." Int J Life Cycle Assess. doi:10.1007/s11367-015-0879-6
[9] Gegick M, Rotella P, Xie T (2010) Identifying Security Bug Reports via Text Mining: An Industrial Case Study. 7th IEEE W Conf Min Softw Repos. doi:10.1109/MSR.2010.5463340
[10] Google (2011) Topic Modeling Tool. https://code.google.com/archive/p/topic-modeling-tool/.
[11] Grubert E (2014) "Life Cycle Topic Modeling Dataset." Stanford University. http://web.stanford.edu/~gruberte/LCA_topics/LCA_abstracts_and_titles_by_year/.
[12] Guinée JB, Heijungs R, Huppes G (2011) Life Cycle Assessment: Past, Present, and Future. Environ Sci Technol 45.1:90–96
[13] Hauschild MZ (2005) Assessing Environmental Impacts in a Life-Cycle Perspective. Environ Sci Technol 39.4:81A–88A
[14] Hertwich EG, Hammitt JK (2001) A Decision-Analytic Framework for Impact Assessment Part I: LCA and Decision Analysis. Int J Life Cycle Assess 6.1:5–12
[15] Hertwich EG, Pease WS, Koshland CP (1997) Evaluating the Environmental Impact of Products and Production Processes: A Comparison of Six Methods. Sci Total Environ 196.1:13–29
[16] International Organization for Standardization (1996) ISO 14000: Environmental Management. ISO, Geneva
[17] International Organization for Standardization (2006) ISO 14044 International Standard. In: Environmental Management–Life Cycle Assessment–Requirements and Guidelines. ISO, Geneva
[18] Jockers M (2011) The LDA Buffet Is Now Open; Or, Latent Dirichlet Allocation for English Majors. Stanford University. https://web.stanford.edu/group/ats/cgi-bin/hivetalkin/?p=2011.
[19] Jockers ML, Witten DM (2010) A Comparative Study of Machine Learning Methods for Authorship Attribution. Lit Linguist Comput. doi:10.1093/llc/fqq001
[20] Johnsen FM, Løkke S (2013) Review of Criteria for Evaluating LCA Weighting Methods. Int J Life Cycle Assess 18.4:840–849
[21] Kostoff RN, Block JA, Solka JL, Briggs MB, Rushenberg RL, Stump JA, Johnson D, Lyons TJ, Wyatt JR (2008) Literature-Related Discovery (LRD): Lessons Learned, and Future Research Directions. Technolog Forecast Social Change 75.2:276–299
[22] Marsmann M, Ryding SO, de Haes HU, Fava J, Owens W, Brady K, Saur K, Schenck R (1999) In Reply to Hertwich & Pease, Int. J. LCA 3 (4) 180 — 181, 'ISO 14042 Restricts Use and Development of Impact Assessment.' Int J Life Cycle Assess 4.2:65–65
[23] Masanet E, Chang Y (2014) Who Cares About Life Cycle Assessment? J Ind Ecol doi:10.1111/jiec.12189

[24] McCallum AK (2002) MALLET: A Machine Learning for Language Toolkit. http://mallet.cs.umass.edu.
[25] McManus MC, Taylor CM (2015) The Changing Nature of Life Cycle Assessment. Biomass Bioenergy. doi:10.1016/j.biombioe.2015.04.024
[26] Miller TR, Neff MW (2013) De-Facto Science Policy in the Making: How Scientists Shape Science Policy and Why It Matters (or, Why STS and STP Scholars Should Socialize). Minerva 51.3:295–315
[27] Nash J, Stoughton MD (1994) Learning to Live with Life Cycle Assessment. Environ Sci Technol 28.5:236A–237A.
[28] Neff MW (2011) What Research Should Be Done and Why? Four Competing Visions among Ecologists. Frontiers Ecol Environ 9.8:462–469
[29] Neff MW, Corley EA (2009) 35 Years and 160,000 Articles: A Bibliometric Exploration of the Evolution of Ecology. Scientometr 80.3:657–682
[30] Owens JW (1997) Life-Cycle Assessment: Constraints on Moving from Inventory to Impact Assessment. J Ind Ecol 1.1:37–49
[31] Plevin RJ, Delucchi MA, Creutzig F (2014) Using Attributional Life Cycle Assessment to Estimate Climate-Change Mitigation Benefits Misleads Policy Makers. J Ind Ecol 18.1:73–83
[32] Salton G (ed) (1971) The SMART Retrieval System: Experiments in Automatic Document Processing. Prentice-Hall. http://jmlr.csail.mit.edu/papers/volume5/lewis04a/a11-smart-stop-list/english.stop.
[33] Scanlon KA, Lloyd SM, Gray GM, Francis RA, LaPuma P (2014) An Approach to Integrating Occupational Safety and Health into Life Cycle Assessment. J Ind Ecol. doi:10.1111/jiec.12146
[34] Schmidt W-P, Sullivan J (2002) Weighting in Life Cycle Assessments in a Global Context. Int J Life Cycle Assess 7.1:5–10
[35] Souza RG, Rosenhead J, Salhofer SP, Valle RAB, Lins MPE (2014) Definition of Sustainability Impact Categories Based on Stakeholder Perspectives. J Clean Prod. doi:10.1016/j.jclepro.2014.09.051
[36] Sullivan D (2001) Document Warehousing and Text Mining: Techniques for Improving Business Operations, Marketing, and Sales. John Wiley & Sons, Inc., New York
[37] Talamini E, Caldarelli CE, Wubben EFM, Dewes H (2012) The Composition and Impact of Stakeholders' Agendas on US Ethanol Production. Energy Policy 50:647–658
[38] Turconi R, Boldrin A, Astrup T (2013) Life Cycle Assessment (LCA) of Electricity Generation Technologies: Overview, Comparability and Limitations. Renew Sust Energy Rev 28:555–565
[39] Vasara P, Rouhiainen J, Lehtinen H (2013) Resource Convergence and Resource Power: Towards New Concepts for Material Efficiency. Philosophical Trans Royal Soc A: Math, Phys Eng Sci 371.1986:20110562
[40] Zamagni A, Amerighi O, Buttol P (2012) Finding Life Cycle Assessment Research Direction with the Aid of Meta-Analysis. J Ind Ecol 16:S39–S52
[41] Zhang Y, Baral A, Bakshi BR (2010) Accounting for Ecosystem Services in Life Cycle Assessment, Part II: Toward an Ecologically Based LCA. Environ Sci Technol 44.7:2624–2631
[42] Zhang Y, Singh S, Bakshi BR (2010) Accounting for Ecosystem Services in Life Cycle Assessment, Part I: A Critical Review. Environ Sci Technol 44.7:2232–2242

Final version available at http://dx.doi.org/10.1007/s11367-016-1153-2.

8 Tables

8.1 Table 1. Three-topic model topic tags and top 20 words

Topic Tag	health	systems analysis	climate change
Word 1	environmental	energy	environmental
Word 2	system	analysis	energy
Word 3	data	cost	production
Word 4	model	emissions	emissions
Word 5	impact	high	impact
Word 6	systems	potential	impacts
Word 7	effects	development	based
Word 8	product	water	system
Word 9	waste	production	waste
Word 10	exposure	power	carbon
Word 11	process	plant	ghg
Word 12	based	generation	systems
Word 13	management	compared	electricity
Word 14	concentrations	important	data
Word 15	population	higher	gas
Word 16	fuel	low	greenhouse
Word 17	methods	related	management
Word 18	toxicity	found	process
Word 19	species	consumption	case
Word 20	factors	lower	products

9 Figure Captions

9.1 Fig. 1 The three-topic model shows a clear tradeoff between topic 1 (labeled "climate change") and topic 2 (labeled "health") between 1995 and today (not all bars add to 100% given the 5% threshold for topic creation)

9.2 Fig. 2 In the four-topic model, a new topic of "waste management" appears and incorporates elements of each of the original 3 topics. Here, waste broadly refers to pollution, which is related to climate change, can cause health issues, and requires good management. The finer topic resolution continues to show the health/climate shift over time (not all bars add to 100% given the 5% threshold for topic creation)

9.3 Fig. 3 In the five-topic model, we see further bifurcation. The association between the risk and health topics is relatively high: words like "exposure" and "toxicity" indicate that the topic labeled "risk assessment" refers at least in part to biological risks. The trend of trading off between health and climate is still visible, though risk assessment (which applies to both themes) remains present in modern literature (not all bars add to 100% given the 5% threshold for topic creation)

9.4 Fig. 4 In the 6-topic model, health and risk assessment remain relatively closely associated. The themes that earlier appear in a "systems analysis" topic start to take on higher specificity as the number of topics increases. With six topics, the health topic and the climate topic do not overlap: the health topic disappears entirely in 2007, the first year the climate topic contributes at all to the abstracts in the sample. The sudden rise of climate-focused literature is evident here, with the climate topic growing to about 25% of recent literature (not all bars add to 100% given the 5% threshold for topic creation)

9.5 Fig. 5 In the 7-topic model, continuing the trend of finer topic gradients, the health, systems, and waste management topics have all split into related but distinct topics. Notably, human health and nonhuman health appear to be distinct. More interesting, the most clearly human health-related topic disappears a few years before the climate topic arrives (not all bars add to 100% given the 5% threshold for topic creation)

10 Figures

10.1 Figure 1. Proportional contribution by topics to abstracts and titles published by year in three-topic model of LCA abstracts indexed in Web of Knowledge, January 1995-May 2014.

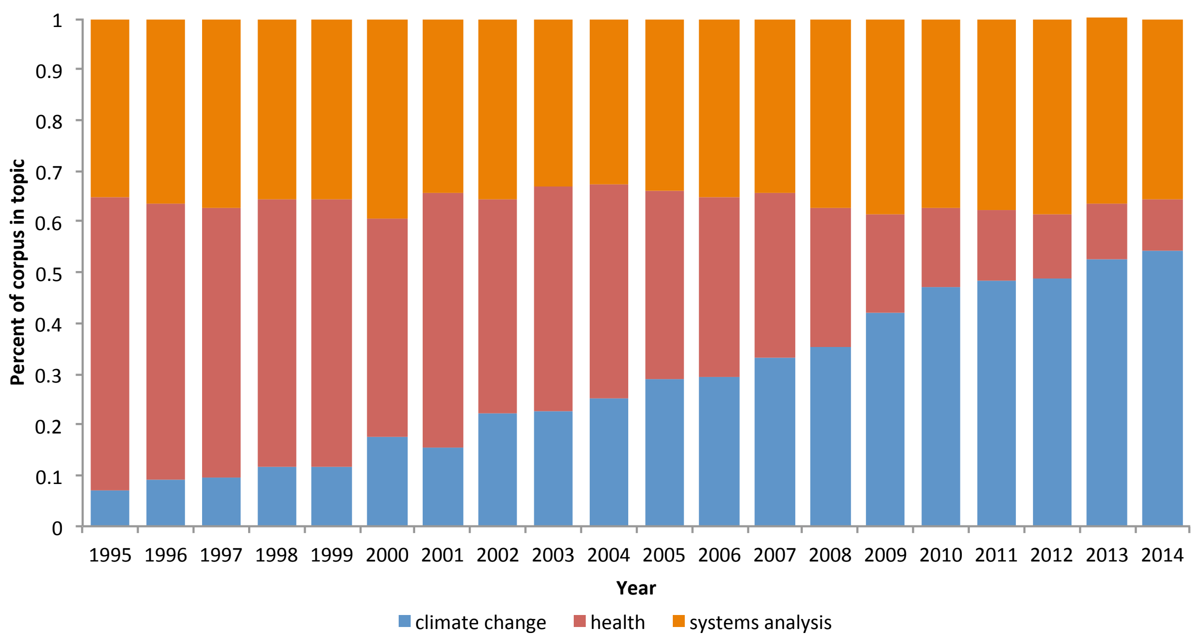

10.2 Figure 2. Proportional contribution by topics to abstracts and titles published by year in four-topic model of LCA abstracts indexed in Web of Knowledge, January 1995-May 2014.

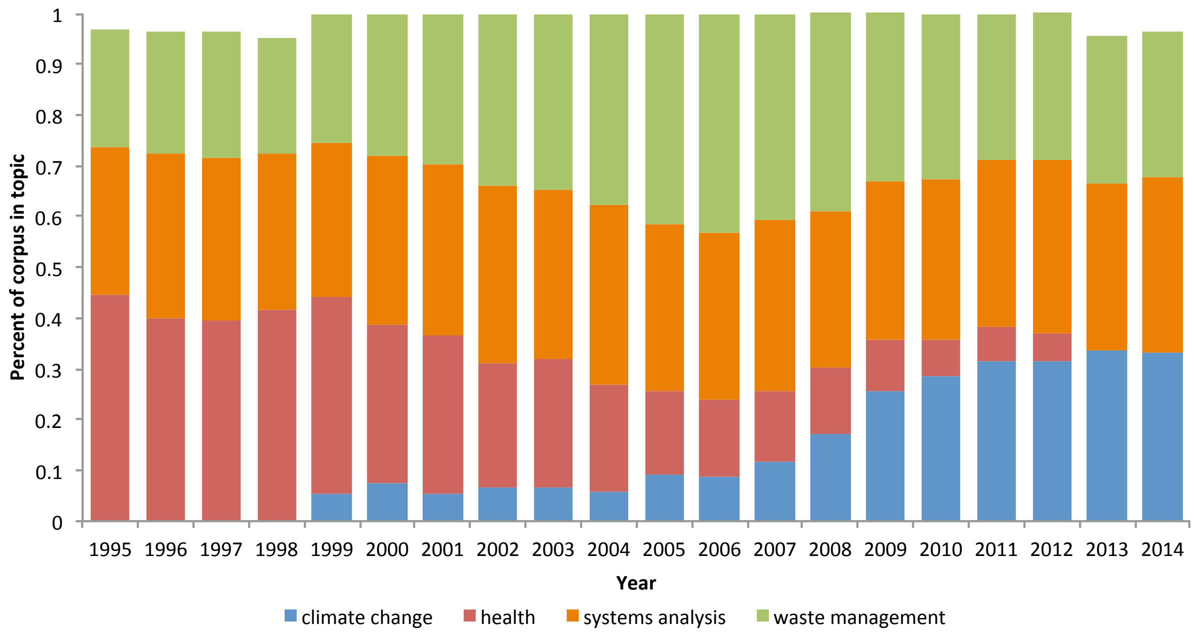

10.3 Figure 3. Proportional contribution by topics to abstracts and titles published by year in five-topic model of LCA abstracts indexed in Web of Knowledge, January 1995-May 2014.

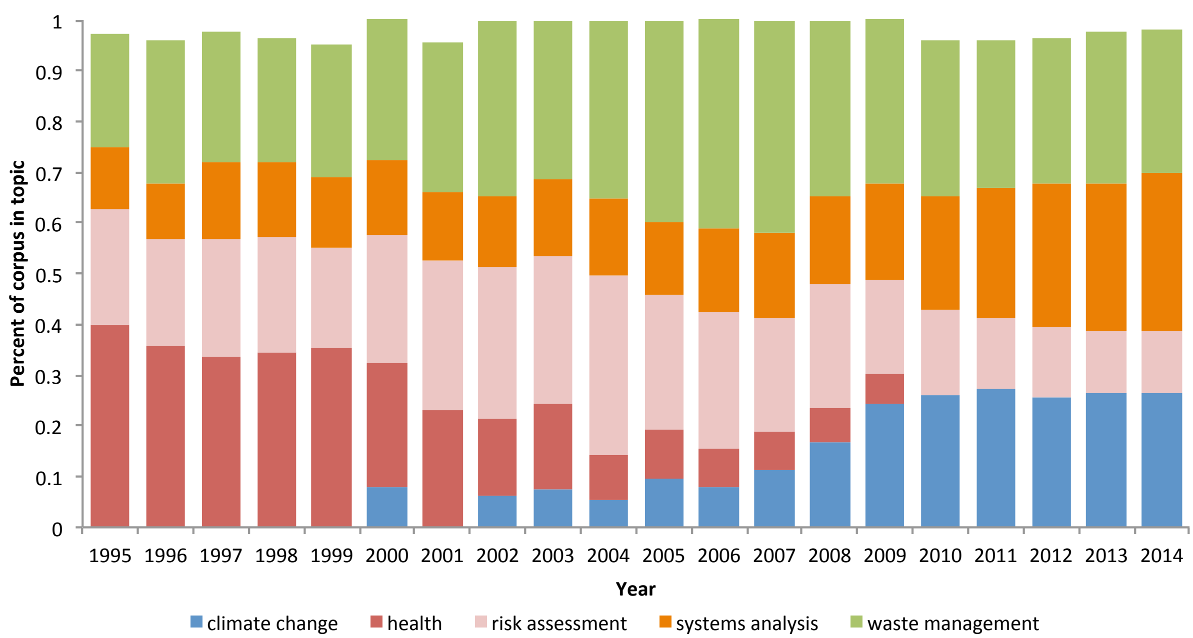

10.4 Figure 4. Proportional contribution by topics to abstracts and titles published by year in six-topic model of LCA abstracts indexed in Web of Knowledge, January 1995-May 2014.

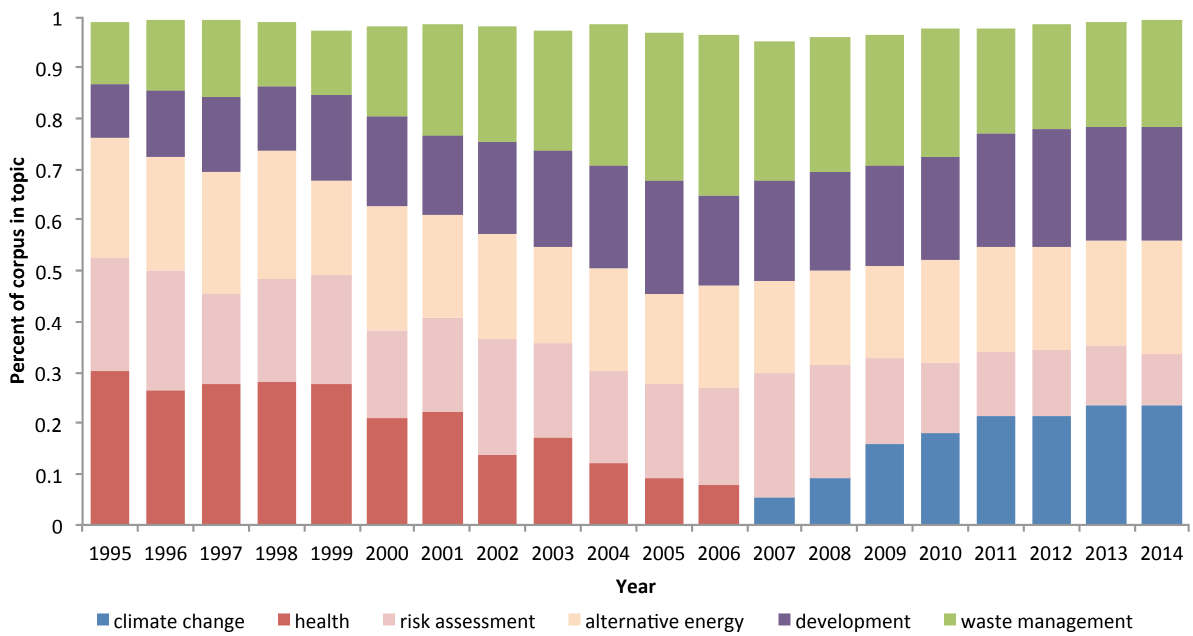

10.5 Figure 5. Proportional contribution by topics to abstracts and titles published by year in seven-topic model of LCA abstracts indexed in Web of Knowledge, January 1995-May 2014.