

 Open access • Posted Content • DOI:10.1101/2019.12.17.880062

## **FoodMine: Exploring Food Contents in Scientific Literature** — [Source link](#)

Forrest Hooton, Giulia Menichetti, Albert-László Barabási, Albert-László Barabási ...+1 more authors

**Institutions:** Northeastern University, Harvard University, Central European University

**Published on:** 18 Dec 2019 - bioRxiv (Cold Spring Harbor Laboratory)

Related papers:

- [Exploring Food Contents in Scientific Literature with FoodMine](#)
- [Food bioactive small molecule databases: Deep boosting for the study of food molecular behaviors](#)
- [Proteomics in food: Quality, safety, microbes, and allergens.](#)
- [Connecting nutrition composition measures to biomedical research](#)
- [The evolution of analytical chemistry methods in foodomics.](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/foodmine-exploring-food-contents-in-scientific-literature-3t23cts8u0>



OPEN

# Exploring food contents in scientific literature with FoodMine

Forrest Hooton<sup>1</sup>, Giulia Menichetti<sup>1</sup> & Albert-László Barabási<sup>1,2,3</sup>✉

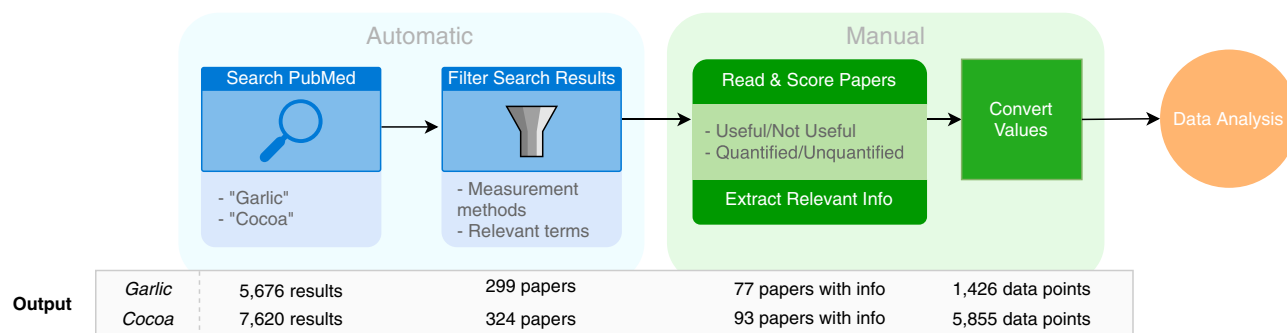
Thanks to the many chemical and nutritional components it carries, diet critically affects human health. However, the currently available comprehensive databases on food composition cover only a tiny fraction of the total number of chemicals present in our food, focusing on the nutritional components essential for our health. Indeed, thousands of other molecules, many of which have well documented health implications, remain untracked. To explore the body of knowledge available on food composition, we built FoodMine, an algorithm that uses natural language processing to identify papers from PubMed that potentially report on the chemical composition of garlic and cocoa. After extracting from each paper information on the reported quantities of chemicals, we find that the scientific literature carries extensive information on the detailed chemical components of food that is currently not integrated in databases. Finally, we use unsupervised machine learning to create chemical embeddings, finding that the chemicals identified by FoodMine tend to have direct health relevance, reflecting the scientific community's focus on health-related chemicals in our food.

Decades of research in nutrition have documented the exceptional role of diet in health, unveiling the role of selected nutrients, like sugars, fats, proteins, vitamins, and other biochemical factors, as well as factors contributing to non-communicable diseases like deficiency diseases, cardiovascular disease, obesity, and diabetes mellitus. However, our ability to explore how food affects our health is severely limited by the lack of systematic knowledge on food composition. Among the national databases for food composition, the most accurate data is maintained by USDA, tracking 188 biochemicals, often called nutritional components<sup>1,2</sup>. Yet, when it comes to the biochemical composition of the food we consume, these nutritional components represent only a tiny fraction of definable biochemicals reported in food. For example, FooDB, a database that integrates food composition data from databases like USDA, Frida, Duke, Phenol Explorer, and others, catalogues altogether 26,625 compounds<sup>3-6</sup>. The majority of these compounds are only identified, with no information on their quantities in specific ingredients. For example, sulfides are reported to be present in the Allium family, like garlic or onion, but the precise quantities for important sulfides like diallyl disulfide (garlic) and dipropenyl sulfide (onion) remain unknown, despite their well-documented role in cancer prevention<sup>7-10</sup>. The current incomplete knowledge of the full biochemical composition of food impedes the research community's ability to uncover the mechanistic effect of the thousands of untracked molecules and their ultimate mechanistic roles in health, achieved either through the microbiome<sup>11</sup>, by contributing to the body's metabolism, or by regulating molecular processes in human cells.

The lack of centralized information on the chemical composition of food does not equal a lack of scientific or commercial interest in these chemicals: an exceptional amount of research focuses on identifying and quantifying the presence of certain chemicals in various foods, as well as the health implications and the biochemical roles of specific food-borne chemicals. The problem is that data on the chemical composition of food is scattered across the multiple research literatures, spanning different scientific communities, from agriculture to food research, and from health sciences to biochemistry. While we witness notable efforts to partially mine this extensive literature and catalogue the scattered data into databases, like Phenol Explorer's focus on polyphenols or eBASIS's prioritization of human intervention studies<sup>6,12,13</sup>, we lack efforts to achieve this across the full food supply and chemicals.

The lack of systematic efforts to map out the existing information on food prompted us to ask how much information is really available on food composition. We developed FoodMine, a pilot project designed to systematically mine the scientific literature to identify and collect all the chemical compositional data for specific ingredients. Hence we demonstrate the capabilities offered by FoodMine by focusing on garlic and cocoa, foods with well documented health effects, which suggests the existence of a sizable yet scattered literature pertaining

<sup>1</sup>Network Science Institute, Northeastern University, Boston, MA, USA. <sup>2</sup>Division of Network Medicine, Department of Medicine, Harvard Medical School, Boston, MA, USA. <sup>3</sup>Department of Network and Data Science, Central European University, Budapest, Hungary. ✉email: alb@neu.edu



**Figure 1.** Overview of data collection process. Starting from PubMed, we retrieved a list of paper titles and abstracts using the Pubmed Entrez API, and then applied text matching to automatically filtered the search results, obtaining a subset of papers, which were then read and manually evaluated. If papers contained information on the chemical content of cocoa or garlic, we manually extracted the relevant information. Finally, we converted values in comparable units. The “Output” bar shows the result of each step for garlic and cocoa.

information on their chemical contents<sup>14,15</sup>. The knowledge gathered here serves as a pilot towards future comprehensive systematic efforts aimed at identifying and organizing the available information on the chemical composition of all food throughout the whole scientific literature.

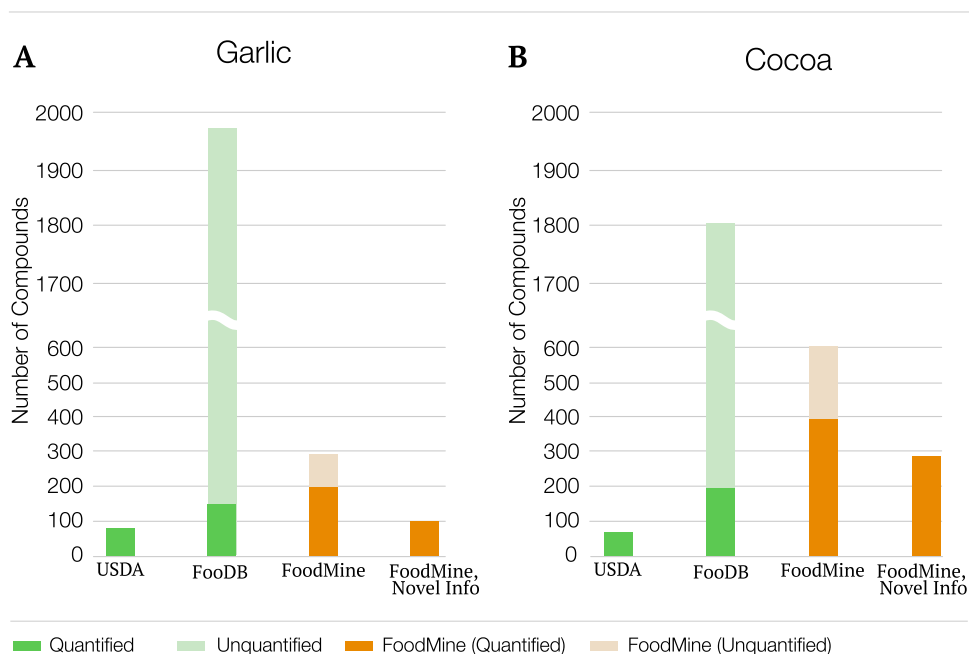
## Results

The FoodMine protocol leveraged the PubMed databases to systematically analyze the title and the abstract of the research papers related to garlic and cocoa (Fig. 1)<sup>16</sup>. We entered each food as a search term, obtaining 5676 papers for garlic and 7620 papers for cocoa. We subset the search results by applying text matching between predefined vocabularies, MeSH terms<sup>17</sup> and the abstract of the paper listed in the PubMed entry, narrowing the results to 415 papers for garlic and 475 papers for cocoa. After obtaining the subset of results, we manually accessed the paper if we could access a “full text link”, downloading 299 papers for garlic and 324 papers for cocoa. Finally, we manually evaluated each paper to identify relevant chemical contents and extract information from it. Of the 623 manually evaluated papers, 77 papers contained chemical composition data for garlic and 93 for cocoa, yielding 1426 and 5855 individual chemical measurements in total for garlic and cocoa, respectively (see Supplementary Material Sect. 1). In the resulting FoodMine database a compound is “quantified” when chemical measurements report absolute contents, and “unquantified” otherwise. Supplementary Fig. S1 shows that the majority of papers for both foods contained only one or at most a few chemical measurements, referred to as records. However, an outlier for cocoa reported 960 records<sup>18</sup>, measuring the contents of several compounds in cocoa sourced from 15 different origins; the permutations of these variables resulted in the high number of data points. Another outlier examined a large spread of compounds related to human taste perception, reporting 68 unique compounds<sup>19</sup>. For garlic, the most relevant outlier reported 198 records<sup>20</sup>.

We integrated the compound records into single compound entries, and manually divided quantified entries into their respective compound class based on FooDB classifications, as shown in Supplementary Fig. S2. We find that ‘Carboxylic Acids and Derivatives’ contains the most explored compounds for both garlic and cocoa, and the ‘Flavonoids’ class is in the top three for both ingredients. Compounds from these two classes are common in plant-based food, hence are expected to be present in garlic and cocoa. We also uncovered reports pertaining to various metallic classes, ‘Toxins’, and ‘Pesticides’. Many compounds in the pesticides class came from a paper focusing on the pesticide residues in cocoa products from local markets in Southwestern Nigeria<sup>21</sup>. Despite its local focus, the examined compounds could directly affect health outcomes worldwide, as Nigeria is the world’s 3rd largest exporter of cocoa<sup>22</sup>.

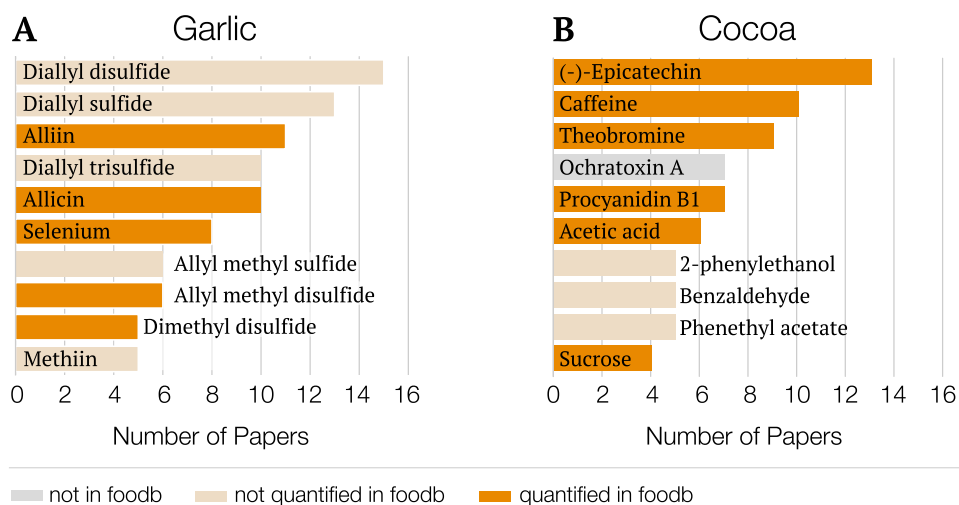
The FooDB and USDA databases allowed us to verify if the information recovered from the literature matches or contrasts the existing knowledge on the composition of these foods (see Supplementary Material Sect. 2 for a detailed description of the comparison methodology). To maximize the coverage of this analysis we merged different variations of garlic and cocoa within the USDA and FooDB databases, like merging “Garlic” and “Soft-necked Garlic” in FooDB when comparing the information to FoodMine. In USDA, all reported compounds are quantified, while FooDB lists both quantified and unquantified compounds. We consider a compound quantified if at least one absolute measurement is reported for the selected foods. Taken together, we find that FoodMine recovered more unique compounds than catalogued by USDA (Fig. 2A,B), and more quantified compounds than catalogued by FooDB. While only 7–9% of compounds are quantified in FooDB and USDA for garlic and cocoa, through FoodMine we collected quantified information for 70% of garlic compounds and 66% of cocoa compounds (see Supplementary Material Sect. 3). For cocoa and garlic, FooDB and USDA contain more unquantified compounds than quantified. However, we find that ~70% of the information reported in the literature was quantified, indicating that the literature contains an extensive body of information currently not recorded in databases (see Supplementary Material Sect. 3). Furthermore, 96 quantified garlic compounds and 283 quantified cocoa compounds are novel, meaning that they were not previously linked to the two ingredients in USDA or FooDB. In summary, 48% and 72% of quantified compounds are novel in both garlic and cocoa, respectively, hence the average increase in quantified measurements offered by FoodMine exceeds 137% (see Supplementary

## Database Compound Comparison



**Figure 2.** Number of unique compounds recovered by FoodMine, USDA, and FooDB. The plots show the number of unique compounds reported by USDA, FooDB, and FoodMine. The columns display (1) the total number of unique quantified compounds in each database, (2) the total number of unique unquantified compounds in each database, and (3) the number of quantified compounds retrieved by FoodMine and never reported before in USDA or FooDB.

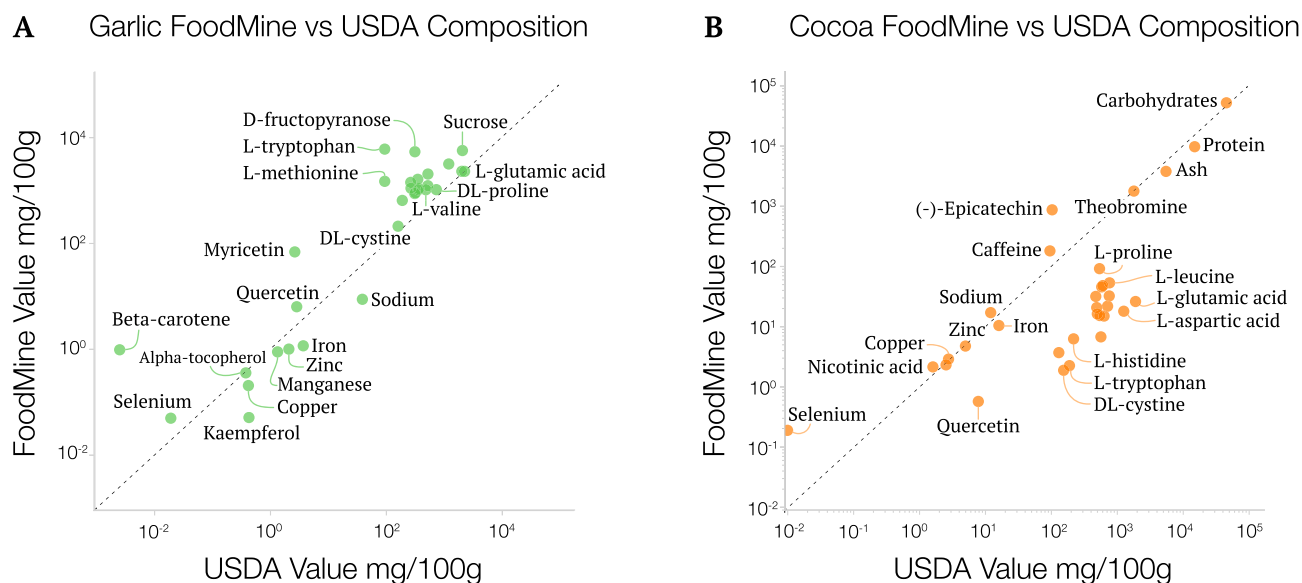
## 10 Most Frequent Compounds



**Figure 3.** Most frequently occurring compounds in FoodMine. The graphs show the top 10 most frequently occurring compounds in terms of number of recovered papers for (A) garlic and (B) cocoa, gauging the research interest in each product. The y-axis displays the compound name, and the x-axis shows the number of papers that contain records for the given compound.

Material Sect. 3). These findings suggest that a systematic mining of the information scattered in the scientific literature could significantly improve our current knowledge of food composition.

The most frequently reported compounds (Fig. 3) in FoodMine are known to play important roles in health effects and flavor. For example, diallyl disulfide is known to contribute to garlic's smell and taste. More importantly, it is implicated in the health effects of garlic, in particular garlic allergy<sup>23,24</sup>. Yet, neither USDA nor FooDB

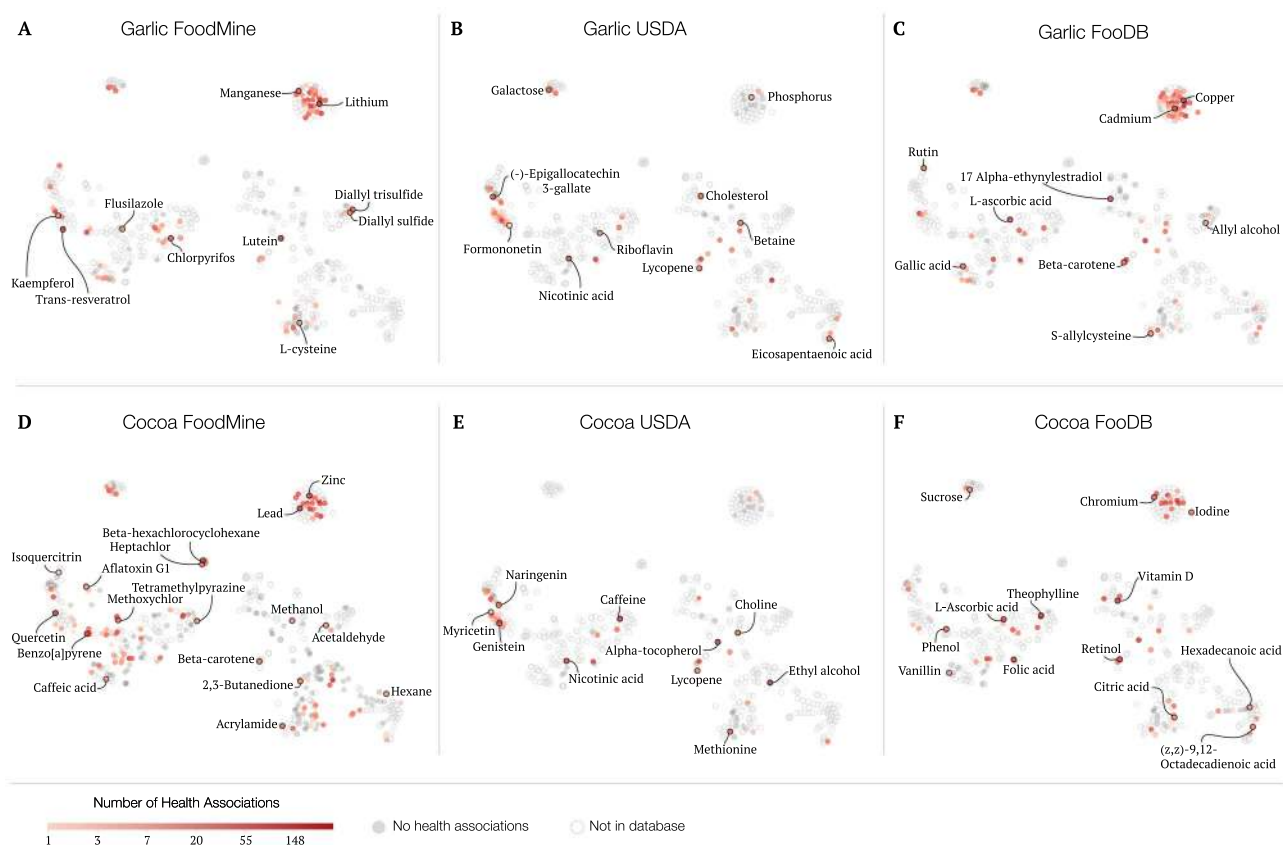


**Figure 4.** Measurement comparisons between FoodMine and USDA. The nutrient concentrations reported by USDA (x-axis), plotted against the content values of matching compounds in FoodMine (y-axis). The dotted line represents the diagonal. We excluded three and two compounds for (A) garlic and (B) cocoa, respectively, because USDA reported zero values for those compounds.

offers quantified information for the compound. This is not an isolated case, as Fig. 3 shows FooDB and USDA lack information on other frequently explored compounds as well. The need to systemically characterize the nutrient profile of a large number of food items, as USDA does, misses information on those compounds that are specific to a few individual foods, despite the potential role they play in health. Indeed, three of the top ten compounds for cocoa are not quantified in FooDB and one is not listed, while for garlic, five of the top ten compounds are not quantified.

To understand the accuracy of the collected quantified data, we compared the FoodMine compound measurements to their corresponding values in USDA, the gold standard for measurement reliability among national food composition databases. Given the limited nutrient panel reported by USDA, we were able to compare only 11% of the chemical compounds we recovered for garlic, and 5% for cocoa. The recovered information spanned a full spectrum of molecules, mixing compounds with both small and large relative quantities (Fig. 4). Overall, we find a good agreement between the FoodMine-recovered and the USDA-reported values (see Supplementary Material Sect. 3 for statistics). Garlic has a logarithmic R-squared value of 0.82, indicating a notable correlation between the known amounts and the FoodMine records, while cocoa only reached 0.56. The lower correlation for cocoa is due to a group of amino acids, reported by papers that examined the contents of roasted cocoa, a processing step that alters the quantities of many chemicals, potentially explaining the difference from the USDA measurements<sup>18,19</sup>. If we remove the data pertaining to roasted cocoa, the logarithmic R-squared increases to 0.75.

To offer a more comprehensive understanding of the classes of chemicals we retrieved and their relation across different databases, we created chemical embeddings using the unsupervised machine learning tool Mol2Vec<sup>25</sup>. Chemical embeddings capture the structural similarity of the chemicals in a low dimensional space. Indeed, as shown in Supplementary Fig. S5, when the chemical classification is known, chemicals belonging to the same class tend to be closer in the embedding space defined by Mol2Vec, suggesting that chemical embeddings successfully capture structural information. This process maps the structural knowledge of the compounds we retrieved, and can be used to integrate further information characterizing the compounds. For instance, given the interest in the association between food-borne chemicals and health outcomes, we can leverage the Comparative Toxicogenomics Database (CTD), that reports manually curated associations between chemicals and diseases<sup>26</sup>. After matching the total number of health implications to each of the quantified chemicals in FoodMine, FooDB, and USDA, we layered this information on the obtained Mol2Vec embedding (Fig. 5). We find that FoodMine covers more chemicals with health effects than FooDB and USDA (see Fig. 5A vs B and C, D vs E and F), a difference particularly clear for cocoa (D vs E and F). Furthermore, we find that the chemicals with health associations are more evenly dispersed throughout the embedding space for FoodMine, implying that FoodMine captures chemicals from rather different chemical classes. Overall, the FoodMine cocoa pilot has detected noticeably more organic, benzenoid, and hydrocarbon compounds, as seen by the absent spaces in E (USDA) and F (FooDB) compared to D (FoodMine) (see Supplementary Material Sect. 3). In summary, compared to the existing databases, FoodMine detects more chemicals with health associations, distributed over a wider range of chemical classes, reflecting a selection bias in the literature: the research community appears to be more focused on chemicals with known health outcomes. Interestingly, there is no overlap between the papers contributing to FoodMine and those manually curated in CTD, meaning that we are recovering information from multiple scientific communities, not only health sciences (see Supplementary Material Sect. 3).



**Figure 5.** TSNE dimensionality reduction of chemical embeddings with health associations. TSNE plots of Mol2Vec chemical embeddings for garlic (A–C) and cocoa (D–F). The colors of each data point encode the number of health implications associated with the compounds based on the CTD database. Dark gray represents chemicals with 0 health associations. We show chemicals catalogued by each studied database for FoodMine (A,D), USDA (B,E), and FooDB (C,F). Markers are filled if the database contains the chemical, and empty if it does not.

## Discussion

Our knowledge pertaining to the more than 26,000 chemicals expected to be present in food, as reported in various databases, is highly incomplete. This incompleteness inspired our efforts to examine how much additional uncatalogued knowledge is scattered in the scientific literature. The invisibility of these compounds to experimental, clinical, epidemiological, and demographic studies—the virtual “dark matter” of nutrients—represents a major roadblock towards a systematic understanding of how diet affects our health. The introduced FoodMine pilot systematically scanned the scientific literature, identifying information about a large number of novel, quantified compounds reported by individual papers. We find that the collected information considerably extends our understanding of food composition. Furthermore, many of the recovered compounds have direct relevance to health and nutrition. For instance, the sulfides, quantified by FoodMine, are responsible for garlic’s unique health effects, yet are currently not quantified in USDA or FooDB.

Garlic and cocoa are only two of the over a thousand natural foods commonly consumed by humans, hence our study supports the hypothesis that there is abundant information in the literature on the composition of other ingredients as well. Indeed, the search terms we used in FoodMine to retrieve papers from PubMed were narrow, and the selection of papers we manually evaluated is small compared to the total body of potential knowledge present in the literature. Consequently, there is likely additional information for garlic and cocoa, not yet captured by FoodMine. Other search terms, focusing on compound classes rather than foods, could uncover an additional body of information about the chemical composition of these ingredients, knowledge that can be generalized to other ingredients as well. For instance, by targeting ‘NEPP’, i.e. non-extractable polyphenols, FoodMine could in principle collect and disambiguate the available literature that reports the food content of this class of chemical compounds, often overlooked by food databases, despite the growing interest on their interaction with the human gut microbiome<sup>27–29</sup>.

Our efforts for garlic and cocoa have proven the existence of a sizable yet scattered literature pertaining to their chemical composition, offering a consistent gain of compositional information compared to what is currently available in food databases. With our pilot we focused on chemical information that has been measured by scientists but was effectively lost to the public, due to the lack of storage and disambiguation in accessible databases. Indeed, despite the complexity characterizing the nutrient dark matter, the food consumption is still far from the efforts of genomic and proteomic research in the construction of biobanks and consortiums, curating

and storing the chemical compounds identified in food. Documenting what is currently known about food composition is a necessary step towards further experimental efforts. In this perspective, the output of FoodMine constitutes a valuable starting point for the creation of standards necessary for targeted metabolomics, helping identify and quantify the variability of these chemical compounds in food<sup>30,31</sup>.

Our next goal is to expand the data collection to multiple staple ingredients. We are prioritizing our search according to the consumption and production statistics available in national and international surveys like NHANES<sup>32</sup> and FAOSTAT<sup>33</sup>, aiming to target foods that would help drastically improve the chemical coverage of our diet, and benefit health studies. While manual curation is still needed to extract measurement details from papers, our machine learning algorithm is ranking the papers in order of relevance, to expedite the data collection. Given the heterogeneous scientific language used to describe food, the second phase of this pilot is key for acquiring additional data training to learn novel language features, such as the occurrence of particular n-grams<sup>34,35</sup>, to maximize the applicability of the algorithm to different foods, without losing precision.

## Methods

Literature mining consisted of three steps: search, selection, and information extraction. We began by searching PubMed with the search term ‘garlic’ and ‘cocoa’, using the Pubmed Entrez Programming Utilities API<sup>36</sup>. After retrieving the PubMed ID’s for search results, we again used the API to retrieve information for each PubMed entry associated with the PubMed ID. We used text matching to scan each PubMed entry’s MeSH terms and abstract for words relevant to biochemicals, food, and pre-selected measurement methodologies, after the API query (see Supplementary Material Sect. 1). The algorithm filtered 5676 results for garlic and 7620 papers for cocoa to 415 and 475 results, respectively. We collected papers from the “Full text links” of the combined 900 entries. We skipped an entry if it did not list any “Full text links” or we did not have access to the paper associated with the links, recovering papers for 299 of the 415 PubMed entries for garlic, and 324 of the 475 entries for cocoa. We also spot-checked search results that fell outside the assessment criteria to quantify the effectiveness of the filtration step. Of those, 0/10 papers contained relevant information for garlic and 3/10 papers contained relevant information for cocoa, indicating that the filtering did eliminate some papers with potentially relevant information.

Papers were individually read by a human assessor to determine whether or not they contained information on the chemical composition of garlic or cocoa. The assessment fell into three categories: “not useful” (papers not containing relevant information for chemical contents in food), “quantified” (papers containing quantitative chemical composition that could be translated to unveil its precise contents in samples), and “unquantified” (paper containing chemical composition, but it could not be converted to unveil precise quantities). Examples of unquantified results were compounds detected by a mass spectrometer that only reported relative percentages, hence we could not record their contents in garlic or cocoa. In addition to the human mining procedure, we used machine learning to create a paper classification algorithm, helping to automatize future data collection (see Supplementary Material Sect. 1). This algorithm takes as input the filtered samples and predicts which papers will contain information on the chemical content of cocoa or garlic. We applied SMOTE sampling technique to balance the labeled data, as only 27% of the reviewed papers contained information on the chemical content of cocoa or garlic<sup>37</sup>. Our algorithm achieved an f1 score of 75.5% on the testing set, better than random. These results are in spite of a limited training set, and could improve with more labels from other foods beyond garlic and cocoa.

All records for a single unique compound were merged into a single entry by calculating the mean of quantified record values. As different papers use different variations of a compound’s name, we applied a chemical disambiguation scheme using PubChem CIDs to add keys to the compounds (see Supplementary Material Sect. 2)<sup>38</sup>. For each entry, we reported the average content value across all data points standardized in units of mg/100 g, and captured additional statistics, such as the highest and lowest reported measurement of the chemical, variance across measurements, and number of measurements. Finally, we leveraged the PubChem CIDs to retrieve a string representation of the structural properties of the molecule (chemical SMILE) which we used as the input for Mol2Vec. Once we learned the vector representation for each chemical, we further reduced the dimensionality using TSNE to obtain the maps shown in Fig. 5 and Supplementary Fig. S5<sup>39</sup>.

## Data availability

The raw data and processing code are available at <https://github.com/fhooton/FoodMine>.

Received: 5 December 2019; Accepted: 9 September 2020

Published online: 01 October 2020

## References

1. USDA. National Nutrient Database for Standard Reference, Release 28 (2015) Documentation and User Guide. **28**, (2015).
2. Bhagwat, S., Haytowitz, D. B. & Holden, J. M. USDA database for the flavonoid content of selected foods release 3. *U.S. Dep. Agric.* 1–156. [https://www.ars.usda.gov/ARUserFiles/80400525/Data/Flav/Flav\\_R03-1.pdf](https://www.ars.usda.gov/ARUserFiles/80400525/Data/Flav/Flav_R03-1.pdf) (2011).
3. FooDB. <https://foodb.ca/>. Accessed 25 June 2019.
4. National Food Institute. Frida Food Data, version 1. *Technical University of Denmark* (2015). <https://frida.fooddata.dk>.
5. U.S. Department of Agriculture, A. R. S. Dr. *Duke’s Phytochemical and Ethnobotanical Databases*. (1992). <https://doi.org/10.15482/USDA.ADC/1239279>.
6. Rothwell, J. A. *et al.* Phenol-Explorer 3.0: A major update of the Phenol-Explorer database to incorporate data on the effects of food processing on polyphenol content. *Database* **2013**, bat070 (2013).
7. Munday, R. & Munday, C. M. Relative activities of organosulfur compounds derived from onions and garlic in increasing tissue activities of quinone reductase and glutathione transferase in rat tissues. *Nutr. Cancer* **40**, 205–210 (2001).
8. Nohara, T. *et al.* Antitumor allium sulfides. *Chem. Pharm. Bull. (Tokyo)* **65**, 209–217 (2017).

9. Wang, H., Yang, J.-H., Hsieh, S.-C. & Sheen, L.-Y. Allyl sulfides inhibit cell growth of skin cancer cells through induction of DNA damage mediated G2/M arrest and apoptosis. *J. Agric. Food Chem.* **58**, 7096–7103 (2010).
10. Nicastro, H. L., Ross, S. A. & Milner, J. A. Garlic and onions: Their cancer prevention properties. *Cancer Prev. Res. (Phila)* **8**, 181–189 (2015).
11. Bashan, A. *et al.* Universality of human microbial dynamics. *Nature* **534**, 259–262 (2016).
12. Kiely, M. *et al.* EuroFIR eBASIS: Application for health claims submissions and evaluations. *Eur. J. Clin. Nutr.* **64**, S101–S107 (2010).
13. Plumb, J. *et al.* eBASIS (Bioactive substances in food information systems) and bioactive intakes: Major updates of the bioactive compound composition and beneficial bioeffects database and the development of a probabilistic model to assess intakes in Europe. *Nutrients* **9**, 1–15 (2017).
14. Garlic and organosulfur compounds. <https://lpi.oregonstate.edu/mic/food-beverages/garlic>.
15. Katz, D. L., Doughty, K. & Ali, A. Cocoa and chocolate in human health and disease. *Antioxid. Redox Signal.* **15**, 2779–2811 (2011).
16. PubMed. <https://pubmed.ncbi.nlm.nih.gov/>.
17. MeSH: Medical Subject Headings. <https://www.nlm.nih.gov/mesh/meshhome.html>.
18. Serra Bonvehí, J. & Ventura Coll, F. Factors affecting the formation of alkylpyrazines during roasting treatment in natural and alkalized cocoa powder. *J. Agric. Food Chem.* **50**, 3743–3750 (2002).
19. Stark, T., Bareuther, S. & Hofmann, T. Molecular definition of the taste of roasted cocoa nibs (*Theobroma cacao*) by means of quantitative studies and sensory experiments. *J. Agric. Food Chem.* **54**, 5530–5539 (2006).
20. Lee, J. & Harnly, J. M. Free amino acid and cysteine sulfoxide composition of 11 garlic (*Allium sativum* L.) cultivars by gas chromatography with flame ionization and mass selective detection. *J. Agric. Food Chem.* **53**, 9100–9104 (2005).
21. Oyekunle, J. A. O., Akindolani, O. A., Sosan, M. B. & Adekunle, A. S. Organochlorine pesticide residues in dried cocoa beans obtained from cocoa stores at Ondo and Ile-Ife, Southwestern Nigeria. *Toxicol. Rep.* **4**, 151–159 (2017).
22. Verter, N. & Bečvářová, V. Analysis of some drivers of cocoa export in Nigeria in the era of trade liberalization. *AGRIS On-line Pap. Econ. Inform.* **06**, 1–11 (2014).
23. Rao, P. *et al.* Diallyl sulfide: Potential use in novel therapeutic interventions in alcohol, drugs, and disease mediated cellular toxicity by targeting cytochrome P450 2E1. *Curr. Drug Metab.* **16**, 486–503 (2015).
24. Garcia-Abujeta, J. L. *et al.* Allergic contact dermatitis to diallyl disulphide in Spain. *J. Allergy Clin. Immunol.* **117**, S130 (2006).
25. Jaeger, S., Fulle, S. & Turk, S. Mol2vec: Unsupervised machine learning approach with chemical intuition. *J. Chem. Inf. Model.* **58**, 27–35 (2018).
26. Davis, A. P. *et al.* The comparative toxicogenomics database: Update 2019. *Nucleic Acids Res.* **47**, D948–D954 (2019).
27. Arranz, S., Silván, J. M. & Saura-Calixto, F. Nonextractable polyphenols, usually ignored, are the major part of dietary polyphenols: A study on the Spanish diet. *Mol. Nutr. Food Res.* **54**, 1646–1658 (2010).
28. Pérez-Jiménez, J., Díaz-Rubio, M. E. & Saura-Calixto, F. Non-extractable polyphenols, a major dietary antioxidant: Occurrence, metabolic fate and health effects. *Nutr. Res. Rev.* **26**, 118–129 (2013).
29. González-Sarrias, A., Espín, J. C. & Tomás-Barberán, F. A. Non-extractable polyphenols produce gut microbiota metabolites that persist in circulation and show anti-inflammatory and free radical-scavenging effects. *Trends Food Sci. Technol.* **69**, 281–288 (2017).
30. Hidalgo, D. *et al.* Biotechnological production of pharmaceuticals and biopharmaceuticals in plant cell and organ cultures. *Curr. Med. Chem.* **25**, 3577–3596 (2018).
31. Moon, K. B. *et al.* Development of systems for the production of plant-derived biopharmaceuticals. *Plants* **9**, 30 (2020).
32. National Health and Nutrition Examination Survey (NHANES). <https://www.cdc.gov/nchs/nhanes/index.htm>.
33. FAOSTAT (Food and agriculture data). <https://www.fao.org/faostat/en/#home>.
34. Jurafsky, D. & Martin, J. H. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition* (Prentice Hall, Upper Saddle River, 2009).
35. Sidorov, G., Velasquez, F., Stamatatos, E., Gelbukh, A. & Chanona-Hernández, L. Syntactic dependency-based n-grams as classification features. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **7630 LNAI**, 1–11 (Springer, Berlin, 2013).
36. Entrez Programming Utilities Help. (2010). <https://www.ncbi.nlm.nih.gov/books/NBK25501/>. Accessed 26 June 2019.
37. Chawla, N. V., Bowyer, K. W., Hall, L. O. & Kegelmeyer, W. P. SMOTE: Synthetic minority over-sampling technique. *J. Artif. Intell. Res.* **16**, 321–357 (2002).
38. PubChem. <https://pubchem.ncbi.nlm.nih.gov/>. Accessed 25 June 2019.
39. Van Der Maaten, L. & Hinton, G. Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**, 2579–2625 (2008).

## Acknowledgements

We thank Daniela Barbery-Flambury for the manual data collection and Dr. Michael Sebek for advice on compound classes.

## Author contributions

F.H. performed data analysis, programming, and contributed to writing the manuscript. G.M. designed the project procedure and analysis, and contributed to writing the manuscript. A.L.B. contributed to interpreting the results and writing the manuscript. G.M. and A.L.B. conceived the project. F.H. and G.M. contributed equally to the project.

## Funding

Funding was provided by American Heart Association (Grant no. 151708), National Institutes of Health (Grant no. 1P01HL132825) and European Research Council (810115-DYNASET).

## Competing interests

ALB is the founder of Scipher Medicine, Foodome, and Nomix companies that leverage the application of big data in health.

## Additional information

**Supplementary information** is available for this paper at <https://doi.org/10.1038/s41598-020-73105-0>.

**Correspondence** and requests for materials should be addressed to A.-L.B.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).



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



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

© The Author(s) 2020