

# 1 **Knowledge Beacons: Web Service Workflow for FAIR** 2 **Data Harvesting of Distributed Biomedical Knowledge**

3 Lance M. Hannestad<sup>1,2</sup>, Vlado Dančik<sup>3</sup>, Meera Godden<sup>1,2</sup>, Wun Ting Suen<sup>1,4</sup>, Kenneth C. Huellas-  
4 Bruskiwicz<sup>1,5</sup>, Benjamin M. Good<sup>6</sup>, Christopher J. Mungall<sup>6¶</sup>, Richard M. Bruskiwicz<sup>1\*¶</sup>

5 <sup>1</sup> STAR Informatics / Delphinai Corporation, Sooke, BC, Canada

6 <sup>2</sup> Department of Computing Science, Simon Fraser University, Burnaby, BC, Canada

7 <sup>3</sup> Chemical Biology and Therapeutics Science Program, The Broad Institute, Cambridge,  
8 MA, United States of America

9 <sup>4</sup> School of Computing Science, University of British Columbia, Vancouver, BC, Canada

10 <sup>5</sup> School of Interactive Arts and Technology, Simon Fraser University, Burnaby, BC, Canada

11 <sup>6</sup> Lawrence Berkeley National Laboratory, Berkeley, CA, United States of America

12 \* Corresponding author

13 E-mail: [richard@starinformatics.com](mailto:richard@starinformatics.com) (RMB)

14 ¶ CJM and RMB are Joint Senior Authors

## 15 **Abstract**

16 The continually expanding distributed global compendium of biomedical knowledge is  
17 diffuse, heterogeneous and huge, posing a serious challenge for biomedical researchers in  
18 knowledge harvesting: accessing, compiling, integrating and interpreting data, information and  
19 knowledge. In order to accelerate research towards effective medical treatments and optimizing

20 health, it is critical that efficient and automated tools for identifying key research concepts and  
21 their experimentally discovered interrelationships are developed.

22 As an activity within the feasibility phase of a project called “Translator”  
23 (<https://ncats.nih.gov/translator>) funded by the National Center for Advancing Translational  
24 Sciences (NCATS) to develop a biomedical science knowledge management platform, we  
25 designed a Representational State Transfer (REST) web services Application Programming  
26 Interface (API) specification, which we call a Knowledge Beacon. Knowledge Beacons provide a  
27 standardized basic workflow for the discovery of concepts, their relationships and associated  
28 supporting evidence from distributed online repositories of biomedical knowledge. This  
29 specification also enforces the annotation of knowledge concepts and statements to the NCATS  
30 endorsed the Biolink Model data model and semantic encoding standards  
31 (<https://biolink.github.io/biolink-model/>). Implementation of this API on top of diverse  
32 knowledge sources potentially enables their uniform integration behind client software which will  
33 facilitate research access and integration of biomedical knowledge.

34 **Availability:** The API and associated software is open source and currently available for  
35 access at <https://github.com/NCATS-Tangerine/translator-knowledge-beacon>.

## 36 Introduction

37 A serious challenge to impactful biomedical research is the one that biomedical researchers  
38 encounter when identifying and accessing pertinent information: the diffuse and voluminous  
39 nature of such data and knowledge. The large, rapidly growing compendium of published scientific  
40 literature is characterized by diverse data encoding standards; numerous, distinct, heterogeneous,  
41 large and often siloed public research data repositories; relatively inaccessible health records;

42 numerous clinical trial and adverse event reports, all spread across disease communities and  
43 biomedical disciplines. The current distributed nature of this knowledge and associated (meta-  
44 )data silos impedes the discovery of related concepts and the relationships between them, an  
45 activity one might call “Knowledge Harvesting”. Many efforts to overcome this challenge focus  
46 on data management principles to make such resources “Findable, Interoperable, Accessible and  
47 Reusable” (FAIR) [1,2].

48 Web access to bioinformatics data spans many generations of web service standards tagged  
49 with many acronyms, e.g. CORBA [3], SOAP/BioMOBY [4] and SADI [5], the latter an exemplar  
50 of the more general paradigm of “Linked Open Data” using OWL/RDF and SPARQL technology,  
51 including Linked Open Fragments [6].

52 A popular web service standard currently in use is the Swagger 2.0 or OpenAPI 3.0  
53 specified REST API (<https://github.com/OAI>). Many extant online biomedical data sources  
54 currently provide such REST API implementations for accessing their data. API registries exist to  
55 index such APIs to facilitate access (for example, the Smart API Registry; <https://smart-api.info/>)  
56 and generalized tools are available to explore the space of such web services (notably, the  
57 Biothings API and Explorer; <https://biothings.io/>). However, the heterogeneity of such APIs can  
58 be a barrier to efficient biomedical knowledge integration.

59 Here we present a REST-based web services specification called the Knowledge Beacon  
60 API (Beacon API) that enables a basic workflow for the discovery of, and navigation through,  
61 biomedical concepts, relationships and associated evidence. This work arises out of an earlier  
62 effort to develop a web application called “*Knowledge.Bio*” [7] to provide enhanced navigation  
63 through the knowledge base of PubMed cited concepts and relationships, captured by text mining

64 in the Semantic Medline Database [8]. The knowledge harvesting workflow underlying  
65 *Knowledge.Bio* is here elaborated into a distributed web service network across diverse knowledge  
66 sources hosted within the NCATS Biomedical Data Translator Consortium, a publicly funded  
67 project supporting the FAIR integration of distributed biomedical research data and knowledge to  
68 accelerate the development of new disease treatments and reduce the barriers between basic  
69 research and clinical advances [9]. The outcome of this work was an iteratively refined web  
70 service specification implemented in an initial set of Beacons, with validation tools and client  
71 applications.

## 72 **Methods**

73 The Knowledge Beacon API is a Swagger 2.0 specification that defines a set of endpoint  
74 paths embodying operations for accessing knowledge sources and discovering shared semantics  
75 for concepts and their relationships (Fig 1).

76 **Fig 1. Knowledge Beacon Workflow.** General step-by-step flowchart illustrating  
77 the sequential invocation of Beacon web service endpoints, with data flows as  
78 indicated. Also enumerated at the bottom left hand corner of the diagram is the set  
79 of metadata endpoints that report semantic terms and namespaces used by the  
80 Beacon in the annotation of results.

81 A Knowledge Beacon (hereafter abbreviated “Beacon”) initiates a workflow for  
82 knowledge discovery by simple search using a concept endpoint either with a *keywords* parameter  
83 (`/concepts?keywords=`) or one with a Compact Uniform Resource Identifier (CURIE;  
84 <https://en.wikipedia.org/wiki/CURIE>) of the concept (`/concepts/{conceptId}`). In both cases, one  
85 or more specific concepts with associated core details are retrieved.

86           Once identified, the canonical CURIE identifier of a chosen concept selected from the  
87 retrieved list is used as an input parameter to access a list of statements about the concept,  
88 documented as subject/predicate/object assertions (**/statements?s=...** where **s** is a subject  
89 canonical concept CURIE). Additional documentation, including supporting citations, associated  
90 with returned statements may be examined by calling the statement’s endpoint again with the  
91 statement identifier of one of the entries returned from the initial call (i.e.  
92 **/statements/{statementId}**).

93           The data model, concept data type (“*category*”) and relationship predicate (“*edge\_label*”,  
94 “*relation*”) terms in results returned by a Beacon are compliant with an emerging public  
95 Biomedical Data Translator Consortium semantic standard and data model, the Biolink Model  
96 (<https://biolink.github.io/biolink-model/>). To assist client data parsing and interpretation, a  
97 Beacon supports several additional endpoints that return metadata summaries of Biolink Model  
98 terms specifically employed by the Beacon to annotate concepts and statements which are  
99 returned: concept type “categories” (**/categories**), identifier name spaces (**/namespaces**),  
100 relationship “predicates” (**/predicates**) plus a “knowledge map” of available subject-predicate-  
101 object triplet statement combinations (**/kmap**).

## 102 **Results**

### 103 **Sample workflow**

104           Knowledge Beacon workflows are implemented as a chained series of REST API endpoint  
105 calls that return data as JSON formatted documents, annotated using Biolink Model standards as  
106 noted above. Here we illustrate a basic minimal two step sequence of such calls which first  
107 identifies a list of concepts with names matching a keyword, then uses the identifier of one returned

108 concept entry to retrieve *subject-predicate-object* statement assertions relating to that selected  
109 concept.

110 **Step 1:** Query knowledge sources by keyword to identify concepts. For example, calling  
111 the basic /concepts endpoint using the Fanconi Anemia complementation group C gene ‘FANCC’  
112 as a keyword, on the Monarch “Biolink API” Beacon, namely:

113 `https://kba.ncats.io/beacon/biolink/concepts?keywords=FANCC`

114 returns the following JSON result with lists of CURIE-identified concepts (one entry shown; full  
115 list of concept entries truncated for conciseness):

```
116     [
117     ... some JSON results
118     {
119         "categories": [
120             "gene",
121             "sequence feature"
122         ],
123         "id": "NCBIGene:102158362",
124         "name": "FANCC"
125     },
126     ... more JSON results
127     ]
```

128 **Step 2:** Using the canonical (URL-encoded) concept CURIE of a selected concept in the  
129 list of concepts returned by keyword in step 1 above, e.g. NCBIGene:102158362, a search for  
130 knowledge assertions (statements) is made on the same database:

131 `https://kba.ncats.io/beacon/biolink/statements?s=NCBIGene%3A102158362`

132

133 This query gives another JSON result which contains asserted “subject-predicate-object”

134 statements about the concept, where the predicate return defines the relationship, as follows:

```
135     [
136     {
137         "id": "biolink:125a0182-0205-44a8-a70a-c03339383177",
138         "object": {
139             "categories": [
140                 "gene"
141             ],
142             "id": "NCBIGene:102158362",
143             "name": "FANCC"
144         },
145         "predicate": {
146             "edge_label": "is_about",
147             "relation": "IAO:0000136"
148         },
149         "subject": {
150             "categories": [
151                 "publication"
152             ],
153             "id": "PMID:17145712",
154             "name": "PMID:17145712"
155         }
156     },
157     ... more JSON results
158 ]
```

159

## 160 **Beacon implementations**

161 A stable set of publicly accessible Beacons are implemented and currently hosted stably  
162 online (as of February 2020) by the NCATS Biomedical Translator Consortium, as enumerated in  
163 Table 1. The Java and Python software implementations of these Beacons are available in  
164 repositories of the NCATS-Tangerine (<https://github.com/NCATS-Tangerine>) GitHub  
165 organization. One implementation is a generic accessor of Biolink Model compliant knowledge  
166 graph databases stored in Neo4j (<https://github.com/NCATS-Tangerine/tkg-beacon>). These  
167 Beacon implementations may be tested using an available validator application  
168 (<https://github.com/NCATS-Tangerine/beacon-validator>). A Python command line Beacon client  
169 is available (<https://github.com/NCATS-Tangerine/tkbeacon-python-client>). A Knowledge  
170 Beacon Aggregator (<https://github.com/NCATS-Tangerine/beacon-aggregator-client>) was also  
171 designed to manage a registered pool of Beacons, and to return consolidated knowledge using  
172 “equivalent concept cliques” to merge related Beacon results.



173 **Table 1. Biomedical Translator Consortium Deployed Beacons**

Subdomain <sup>a</sup>	Beacon Description	Wrapped Knowledge Source
semmeddb	Semantic Medline Database [8]	<a href="https://skr3.nlm.nih.gov/SemMedDB/">https://skr3.nlm.nih.gov/SemMedDB/</a>
biolink	Monarch Database Biolink API [10]	<a href="https://api.monarchinitiative.org/api/">https://api.monarchinitiative.org/api/</a>
hmdb	Human Metabolome Database [11]	<a href="http://www.hmdb.ca/">http://www.hmdb.ca/</a>
rhea	Rhea Annotated Biochemical Reactions database [12]	<a href="https://www.rhea-db.org/">https://www.rhea-db.org/</a>
smpdb	Small Molecular Pathway Database [13]	<a href="http://smpdb.ca/">http://smpdb.ca/</a>
ndex	nDex Bio Graph Archive [14]	<a href="http://www.ndexbio.org">http://www.ndexbio.org</a>

174 <sup>a</sup> The basepath of each Beacon has the form **<https://kba.ncats.io/beacon/<Subdomain>>**, where  
175 the **<Subdomain>** is as listed in column 1 of the table.

## 176 Discussion

177 The Knowledge Beacon API is a basic knowledge discovery workflow (Figure 1)  
178 representing a relatively high-level use case of user interaction with the biomedical knowledge  
179 space, and as such, lacks the full expressive power of a general knowledge query language  
180 interface like SPARQL. Furthermore, the beacon data model aligns with the emerging Biolink  
181 Model standards of the Biomedical Translator Consortium as its template for knowledge  
182 representation. As such, Beacons do not automatically express results in a generic manner as do  
183 knowledge representations such as RDF, although conversion of Beacon statement results into  
184 RDF format is easily accomplished. Finally, aside from some general profiling of the performance  
185 of the Swagger API endpoints on various knowledge sources, we have not here conducted a  
186 rigorous computing-theoretic assessment of the efficiency of this form of knowledge harvesting,

187 although early experience with Beacons point to challenges with internet latency and knowledge-  
188 source specific differences in query performance. In  
189 prototyped a “Knowledge Beacon Aggregator” to provide enhanced asynchronous  
190 query/status/retrieval endpoints as a client-friendly integration layer for managing access to, and  
191 merging data from, a registered catalog of multiple Beacon implementations.

192 Despite the use of some off-the-shelf API generation tools, the wrapping of knowledge  
193 sources as Beacons remains a labour-intensive activity. The semantics of the knowledge source  
194 being wrapped must be heuristically translated. This is somewhat easier for knowledge sources  
195 which have a small number of easily resolved discrete data types (i.e. discrete Biolink Model  
196 concept categories of data) and namespaces with clear mapping onto those discrete data types.

197 In contrast, some “graph” knowledge sources, for example, the NDex Bio biomedical  
198 network data archive (<https://home.ndexbio.org/index/>, wrapped by this project as the *ndex*  
199 Beacon), don’t have such clear concept category and relationship predicate tagging of much of the  
200 archived data. The development of useful but (so far) imprecise heuristics to tag such data on the  
201 fly is required to develop a useful Beacon. In other cases, such as biomedical knowledge resources  
202 whose data object namespace aggregates several types of concepts in a fuzzy manner with limited  
203 additional concept category tagging, it may be even more challenging to semantically tag data  
204 entries for beacon export.

205 A few common library and reference implementations are developed for Beacons; the  
206 Beacon platform would benefit from the further development of standardized tools to  
207 systematically assist such wrapping of native knowledge sources.

208           The availability of a shared API standard for knowledge integration doesn't, in and of itself,  
209 deal with all the challenges of FAIR data integration within the global community. Practical  
210 experience with knowledge harvesting using such API implementations has revealed performance  
211 issues relating to internet and service latency, bandwidth limitations. Knowledge warehousing in  
212 centralized knowledge graphs using ETL (Extract, Transform, Load) processes may sometimes  
213 result in a more tractable process for biomedical knowledge integration; however, such approaches  
214 are still faced with the task of merging equivalent concepts, including the elimination of duplicate  
215 concepts and the resolution of conflicting information, including weighting of assertions differing  
216 in levels of confidence. More unique to ETL warehousing approaches is the ongoing problem of  
217 keeping such resources up-to-date relative to their original knowledge sources. Note that ETL  
218 warehouses and API driven distributed knowledge harvesting approaches can be complementary,  
219 in that ETL data warehouses can also themselves be accessed by the application of web service  
220 REST API's like the Knowledge Beacon API. In fact, some of the current Beacon implementations  
221 use this approach: a back end Biolink Model compliant Neo4j knowledge graph directly wrapped  
222 with the API.

223           The Linked Open Data paradigm using RDF knowledge representation and SPARQL  
224 represents an alternate paradigm for distributed knowledge integration, the theoretical performance  
225 of which was surveyed by Verborgh *et al* [6]. In their assessment, it was noted that downloadable  
226 RDF knowledge data sets and SPARQL endpoints to triple store knowledge bases represent two  
227 extremes of a continuum of RDF knowledge access, each with their characteristic advantages and  
228 weaknesses. They proposed that a constrained query selector specification and RDF representation  
229 – with data, metadata and hypermedia controls - denoted as Linked Data Fragments could be a  
230 shared design representation spanning both ends of the continuum. Furthermore, they proposed an

231 intermediate implementation - termed Triple Pattern Fragments - partitioning RDF processing  
232 more symmetrically across client and server, thus potentially mitigate some of the challenges of  
233 both ends of the API design continuum, for more balanced client-server performance and greater  
234 ease of implementation (see <https://github.com/LinkedDataFragments>).

235 Generally, API approaches to knowledge harvesting may work best with use cases  
236 involving smaller batches of knowledge retrieval based on a focused navigation of the knowledge  
237 space from larger open-ended data sources which would be refractory to import into centralized  
238 knowledge graphs.

239 Finally, there are two other API standards of the Biomedical Data Translator Consortium:  
240 the “NCATS Reasoner API” (<https://github.com/NCATS-Tangerine/NCATS-ReasonerStdAPI>)  
241 and the Biothings API (<https://biothings.io/>). Although there are parallels between them, the  
242 Beacon API is a simpler lower level interface to knowledge resources than the Reasoner API, and  
243 is somewhat more constrained to the Biolink Model than the Biothings API. But the utility of the  
244 Reasoner API is inspiring efforts to publish a Reasoner API interface on top of an implementation  
245 of Knowledge Beacons (<https://github.com/NCATS-Tangerine/kba-reasoner>). We did also  
246 implement prototype Beacon wrapper for the Biothings API ([https://github.com/NCATS-](https://github.com/NCATS-Tangerine/biothings-explorer-beacon)  
247 [Tangerine/biothings-explorer-beacon](https://github.com/NCATS-Tangerine/biothings-explorer-beacon)).

## 248 **Availability**

249 Knowledge Beacon software is open source licensed and available for access in GitHub. A  
250 suitable introduction to the API, containing references to related software components, can be  
251 found at <https://github.com/NCATS-Tangerine/translator-knowledge-beacon>.

## 252 Acknowledgments

253 BMG and RMB collaborated on the predecessor “Knowledge.Bio” application embodying  
254 the workflow captured by the Knowledge Beacon API. CJM and BMG coined the name  
255 “Knowledge Beacon” to express the architectural vision of uniformly wrapped knowledge sources  
256 for distributed knowledge discovery and harvesting. RMB and his team elaborated the original  
257 software design of web service endpoints and initial code implementations embodying Beacon  
258 functionality, then guided further iterations of the API based on feedback from colleagues within  
259 the Biomedical Translator Consortium, with special mention to co-author VD who proposed  
260 insightful revisions to the API, based on his direct experience implementing a Beacon to wrap  
261 HMDB. Under overall supervision by RMB, the heavy lifting of iterative software development  
262 of Beacon implementations - including several beacons, client (including aggregator) and  
263 validation applications - was undertaken by LMH while he was a member of the STAR Informatics  
264 team, assisted by the valuable software programming contributions of several computing science  
265 cooperative education students: MG, WTS and KCHB.

266 The authors would like to sincerely thank Nomi Harris and Marcin Joachimiak of LBNL  
267 for their very helpful editorial feedback on, and suggested revisions to the draft manuscript.

268 The authors would also like to thank the various members of the Biomedical Data  
269 Translator Consortium who gave helpful user needs feedback and support of the Knowledge  
270 Beacon API during its development, in particular, Chris Bizon and Stephen Ramsay. We also  
271 acknowledge here Greg Stupp who, while employed at TSRI, implemented an earlier version of a  
272 Beacon wrapper for biomedical knowledge in Wikidata.

## 273 **References**

- 274 1. Wilkinson MD, Dumontier M, Aalbersberg IJ, Appleton G, Axton M, Baak A, et al. The  
275 FAIR Guiding Principles for scientific data management and stewardship. *Sci Data* 2016  
276 Mar 16;3:160018. doi:10.1038/sdata.2016.18. PubMed PMID: 26978244; PubMed  
277 Central PMCID: PMC4792175
- 278 2. Wilkinson MD, Dumontier M, Jan Aalbersberg I, Appleton G, Axton M, Baak A, et al.  
279 Addendum: The FAIR Guiding Principles for scientific data management and stewardship.  
280 *Sci Data* 2019 Mar 19;6(1):6. doi: 10.1038/s41597-019-0009-6. PubMed PMID:  
281 30890711; PubMed Central PMCID: PMC6427092.
- 282 3. Wilkinson MD, McCarthy L, Vandervalk B, Withers D, Kawas E, Samadian S. SADI,  
283 SHARE, and the *in silico* scientific method *BMC Bioinformatics*. 2010 Dec 21;11 Suppl  
284 12:S7. doi: 10.1186/1471-2105-11-S12-S7. PubMed PMID: 21210986; PubMed Central  
285 PMCID: PMC3040533.
- 286 4. BioMoby Consortium, Wilkinson MD, Senger M, Kawas E, Bruskiwich R, Gouzy J, et  
287 al. Interoperability with Moby 1.0--it's better than sharing your toothbrush! *Brief*  
288 *Bioinform*. 2008 Jan 31;9(3):220-31. doi:10.1093/bib/bbn003. PubMed PMID: 18238804.
- 289 5. Stevens R, Miller C. (2000) Wrapping and interoperating bioinformatics resources using  
290 CORBA. *Brief Bioinform*. 2000 Feb;1(1):9-21. PubMed PMID: 11466976.
- 291 6. Verborgh R, Vander Sande M, Hartig O, Van Herwegen J, De Vocht L, De Meester B, et  
292 al. Triple Pattern Fragments: a Low-cost Knowledge Graph Interface for the Web. *Journal*  
293 *of Web Semantics*. 2016 Mar; 37:184-206. doi:10.1016/j.websem.2016.03.003

- 294 7. Bruskiwich RM, Huellas-Bruskiwicz KC, Ahmed F, Kaliyaperumal R, Thompson M,  
295 Erik Schultes E, et al. Knowledge.Bio: A web application for exploring, building and  
296 sharing webs of biomedical relationships mined from PubMed. 2016.  
297 doi: <http://dx.doi.org/10.1101/055525>
- 298 8. Kilicoglu H, Shin D, Fiszman M, Rosemlat G, Rindflesch TC. SemMedDB: a PubMed-  
299 scale repository of biomedical semantic predications. *Bioinformatics*. 2012 Dec  
300 1;28(23):3158-60. doi: 10.1093/bioinformatics/bts591. PubMed PMID: 23044550;  
301 PubMed Central PMCID: PMC3509487.
- 302 9. The Biomedical Data Translator Consortium. Toward A Universal Biomedical Data  
303 Translator. *Clin Transl Sci*. 2019 Mar;12(2):86-90. doi: 10.1111/cts.12591. . PubMed  
304 PMID: 30412337; PubMed Central PMCID: PMC6440568.
- 305 10. Shefchek KA, Harris NL, Gargano M, Matentzoglou N, Unni D, Brush M, et al. The  
306 Monarch Initiative in 2019: an integrative data and analytic platform connecting  
307 phenotypes to genotypes across species. *Nucleic Acids Res*. 2020 Jan 8;48(D1):D704-  
308 D715. doi: 10.1093/nar/gkz997. PubMed PMID: 31701156.
- 309 11. Wishart DS, Feunang YD, Marcu A, Guo AC, Liang K, Vázquez-Fresno R, et al. HMDB  
310 4.0: the human metabolome database for 2018. *Nucleic Acids Res*. 2018 Jan  
311 4;46(D1):D608-D617. doi: 10.1093/nar/gkx1089. PubMed PMID: 29140435; PubMed  
312 Central PMCID: PMC5753273.
- 313 12. Morgat A, Lombardot T, Axelsen KB, Aimo L, Niknejad A, Hyka-Nouspikel N, et al.  
314 Updates in Rhea - an expert curated resource of biochemical reactions. *Nucleic Acids Res*.

- 315           2017 Apr 20;45(7):4279. doi: 10.1093/nar/gkw1299. PubMed PMID: 27980062; PubMed  
316           Central PMCID: PMC5397180.
- 317       13. Jewison T, Su Y, Disfany FM, Liang Y, Knox C, Maciejewski A, et al. SMPDB 2.0: big  
318           improvements to the Small Molecule Pathway Database. *Nucleic Acids Res.* 2014  
319           Jan;42(Database issue):D478-84. doi: 10.1093/nar/gkt1067. Epub 2013 Nov 6. PubMed  
320           PMID: 24203708; PubMed Central PMCID: PMC3965088.
- 321       14. Pratt D, Chen J, Pillich R, Rynkov V, Gary A, Demchak B, Ideker T. NDEX 2.0: A  
322           Clearinghouse for Research on Cancer Pathways. *Cancer Res.* 2017 Nov 1;77(21):e58-e61.  
323           doi: 10.1158/0008-5472.CAN-17-0606. PubMed PMID: 29092941; PubMed Central  
324           PMCID: PMC5679399.



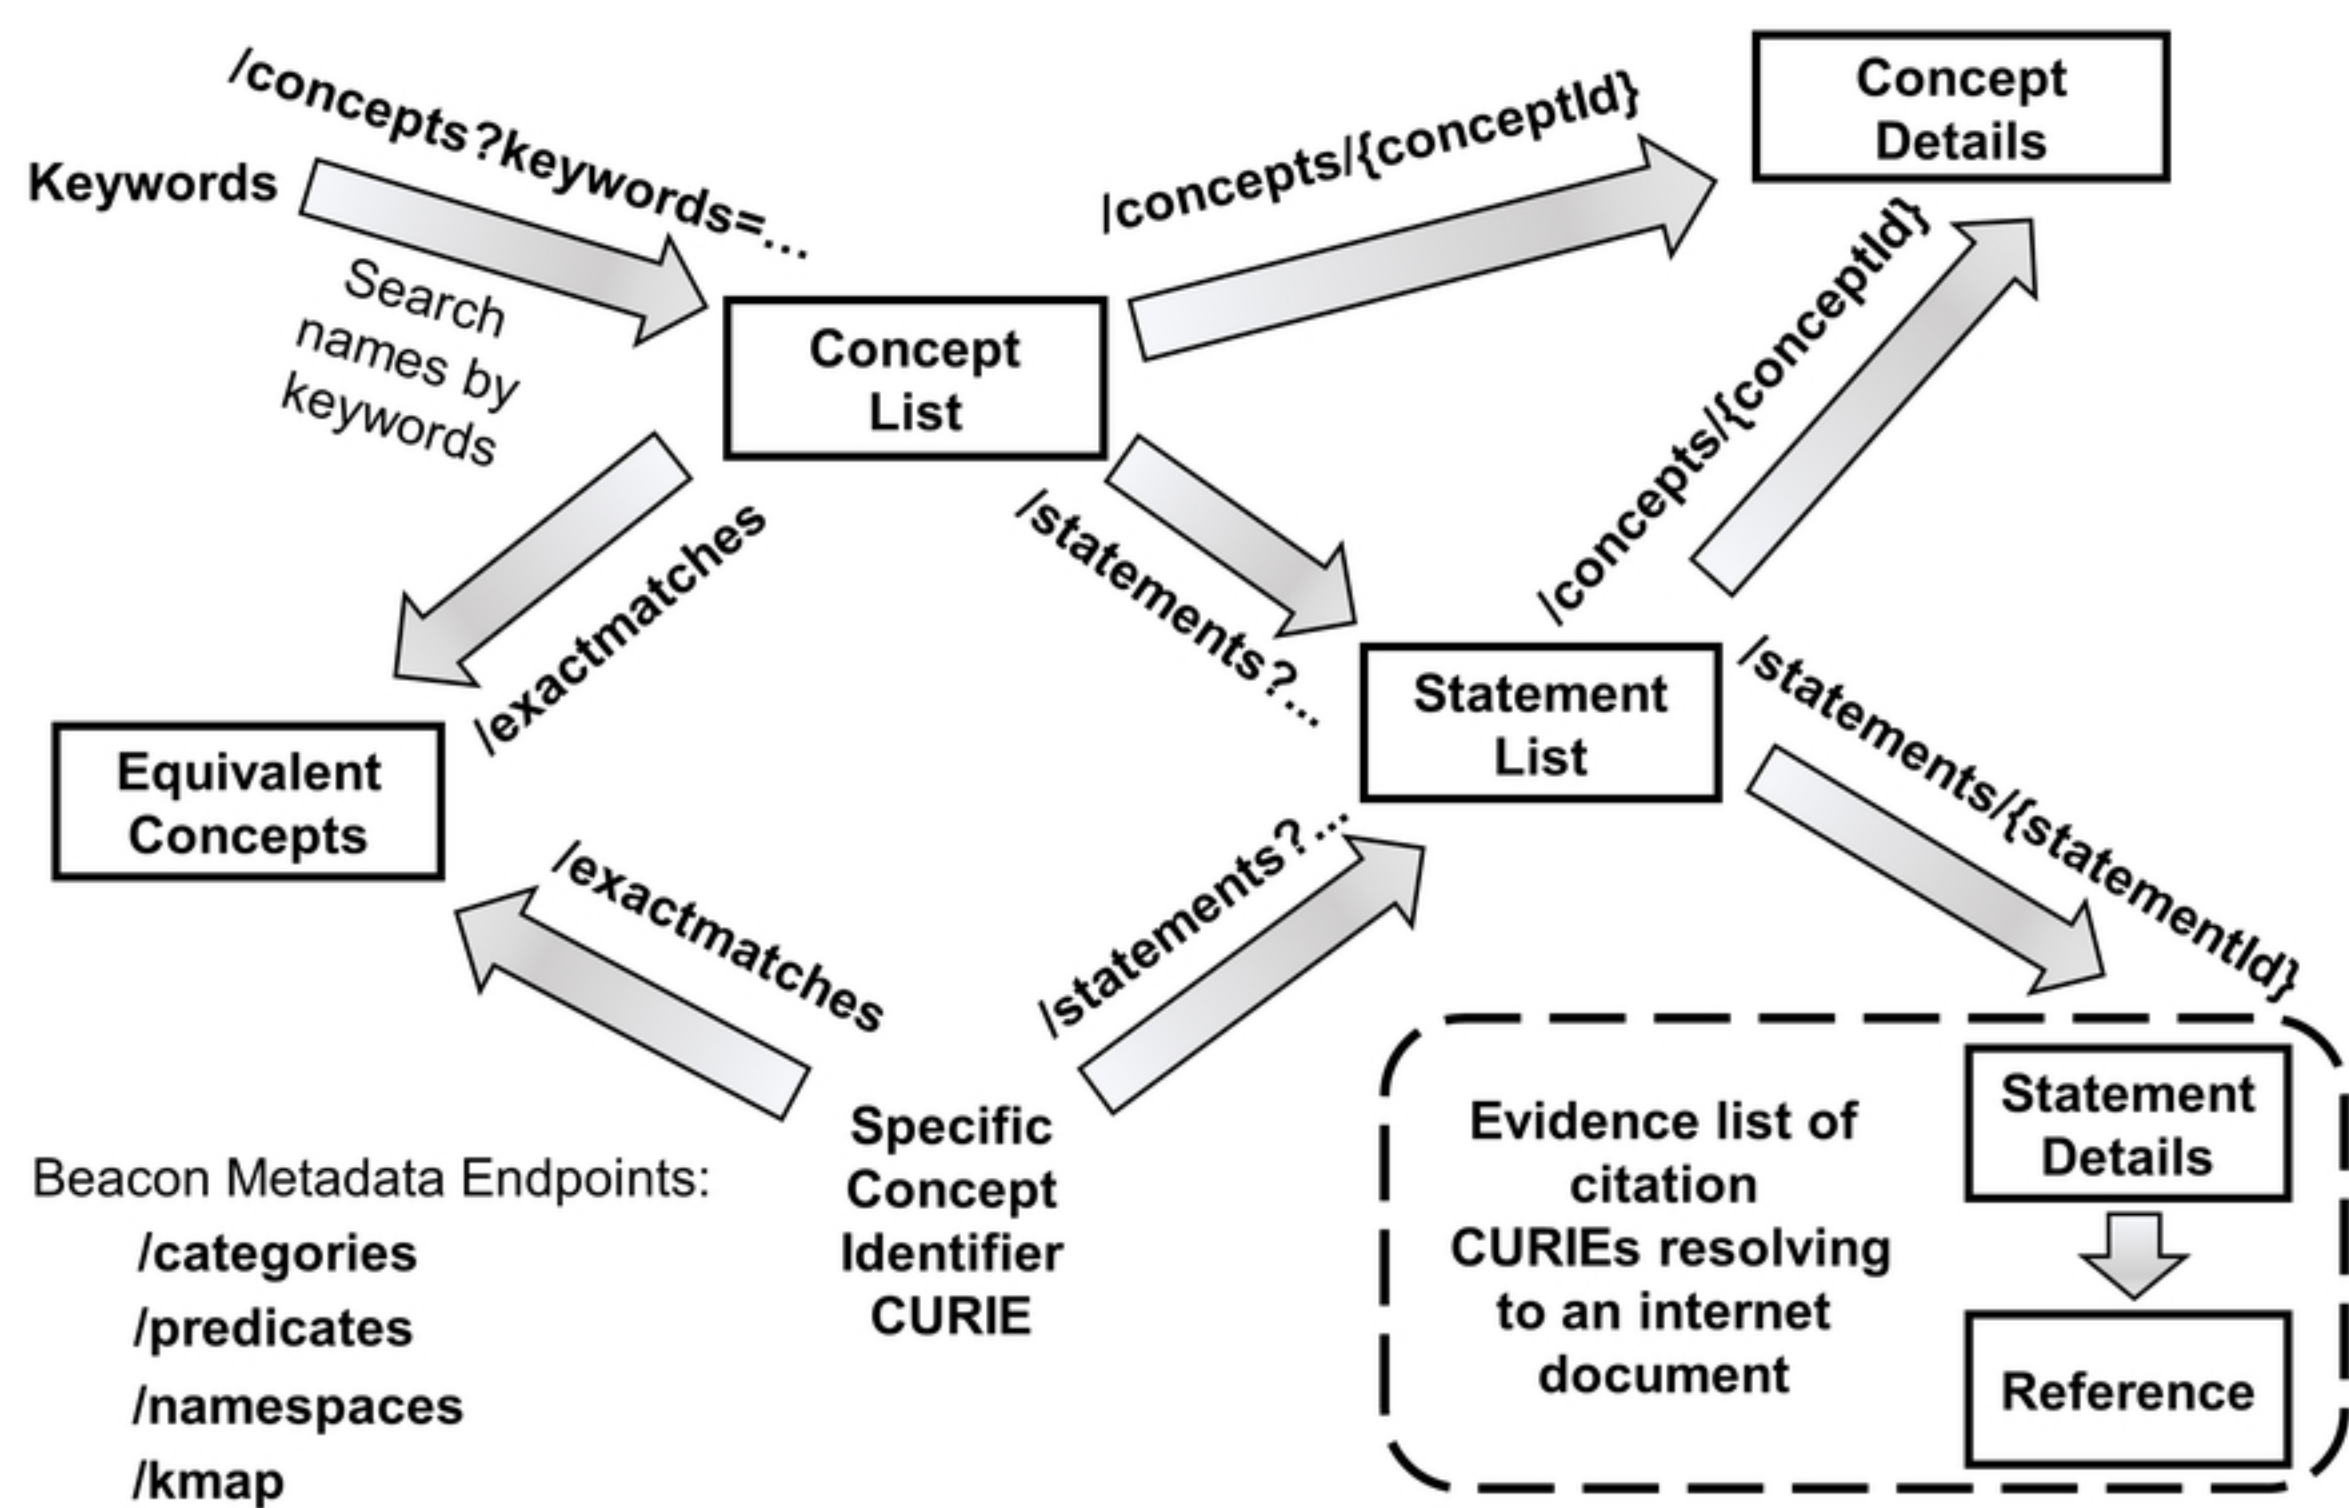


Figure 1