
The Largest Knowledge Graph in Materials Science - Entities, Relations, and Link Prediction through Graph Representation Learning

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 This paper introduces MatKG, a novel graph database of key concepts in material
2 science spanning the traditional material-structure-property-processing paradigm.
3 MatKG is autonomously generated through transformer-based, large language
4 models and generates pseudo ontological schema through statistical co-occurrence
5 mapping. At present, MatKG contains over 2 million unique relationship triples
6 derived from 80,000 entities. This allows the curated analysis, querying, and
7 visualization of materials knowledge at unique resolution and scale. Further,
8 Knowledge Graph Embedding models are used to learn embedding representations
9 of nodes in the graph which are used for downstream tasks such as link prediction
10 and entity disambiguation. MatKG allows the rapid dissemination and assimilation
11 of data when used as a knowledge base, while enabling the discovery of new
12 relations when trained as an embedding model.

13 1 Introduction

14 Comprehensive knowledge of a given material requires the integration of disparate streams of informa-
15 tion that include compositional data, thermodynamic parameters, applications, phase/symmetry labels,
16 synthesis and processing routines, as well as physical, chemical, thermal, optical, and functional
17 properties[1]. In general, it is difficult to find all this information in one place, with the result that
18 comprehensive knowledge of a given material is often missing, even when the data are available.
19 Given the rate at which new data are being accumulated, the amount of available data is far greater
20 than what can be accessed or assimilated. The standard paradigm of data sharing and storage - through
21 peer reviewed scientific publications and relational databases - remains inadequate for the Materials
22 Genome Age where artificial intelligence is increasingly employed to accelerate materials discovery
23 and design [2, 3, 4, 5]. The task of data organization has been approached through custom ontologies
24 that build relations between data points through manual expert input. While several domain specific
25 ontologies such as Nanomine[6], Chemos[7], etc. have been written over the years, no field-wide
26 ontology exists focused on materials science. Given the onerous task of assigning a relation among
27 individual pairs of data, even highly generalizable ontologies such as SKOS[8] have not been applied
28 to materials at scale.

29 In this paper, we introduce MatKG, a novel graph database that links major conceptual entities in the
30 discipline using transformer-based large language models. The database is autonomously extracted
31 from over 4 million papers on the topic of materials and includes chemistry, structure, property,
32 application, synthesis, and characterization data that are aggregated in the form of relational triples
33 <subject, predicate, object>. MatKG has over 2 million unique relationships among over 80,000
34 unique entities.

35 2 Methods

36 **Entity Generation:** A Named Entity Recognition (NER) [9] model was used to extract 80,000
37 unique entities from the abstracts and figure captions of over 4 million scientific publications [10]
38 in the field of material science. Being information dense, these contain low ‘noise’ and are hence
39 particularly suitable for large scale autonomous data mining[11, 12, 13]. The NER model follows
40 the scheme developed in MatScholar [14] and is built on MatBERT [14], a Large Language Model
41 (LLM) trained on a material science text corpus that classifies text tokens into one of the following
42 seven categories: Material (CHM), Property (PRO), Application (APL), Synthesis Method (SYN),
43 Characterization Method (CMT), Descriptor (DSC), and Symmetry/Phase Label (SPL). Derived
44 from the traditional structure-property-processing-application paradigm in material science[1], these
45 entities encapsulate the sum total of the knowledge of any given concept, be it a particular chemistry,
46 process, property, or application. Where possible, each entity is linked to an identifier in Wikipedia
47 using procedure developed in [15] or the corresponding descriptor page in the Materials Project[16].
48 This allows the mapping of entities to broader knowledge bases such as DBpedia[17] and YAGO[18],
49 thereby allowing holistic integration of MatKG with the larger knowledge graph community.

50 **Link Generation :** If entities e_1 and e_2 have the NER tags $T[e_1]$ and $T[e_2]$, they are assigned the
51 relationship $T[e_1]_T[e_2]$ and the weight $v(e_1; e_2)$ according to the method detailed in Appendix 5.1.
52 Subsequently, they are either filtered based on a predefined threshold to form knowledge triples of
53 the form $\langle e_1, T[e_1]_T[e_2], e_2 \rangle$ (1) or as a quartet of the form $\langle e_1, T[e_1]_T[e_2], e_2, v(e_1; e_2) \rangle$ (2)
54 (See Appendix 5.1). (1) allows the extraction of 160,000 high fidelity links between about 12,000
55 unique entities, while (2) results in 2 million relations from up to 80,000 unique entities, thereby
56 demonstrating that a weighted link extraction approach captures far more data - increasing the scope
57 of the knowledge base.

58 **Graph Representation Learning :** The vector representations for the entities in the graph are
59 learnt using knowledge graph embedding models (KGE)[19],[20], [21]. The models are evaluated
60 using mean reciprocal rank (MRR) and hits@(1,10,100) metrics on the test set as described in KGE
61 literature [22]. All models are implemented using the publicly available AmpliGraph Library[23].
62 The model with the highest MRR on the test set was used to perform downstream tasks that are
63 described later.

64 3 Results

65 3.1 Knowledge base creation

66 The autonomously created highly interconnected knowledge graph for materials consists of the seven
67 NER categories and 49 relations (including inverse relations such as *APL_PRO* and *PRO_APL*).
68 The KG is thus a bidirectional digraph. The three most common types of entities are *PRO*, *CHM*,
69 and *CMT*, while the most frequent relations are *CHM_PRO*, *PRO_DSC*, and *CHM_CMT*
70 (see Appendix, Table 1, 2). The large number of material-property (108 k) and material-application
71 (89 k) triples could correspond to the type of information usually present in abstracts, while char-
72 acterization related information originate from figure captions. Many papers in the corpus relate to
73 inorganic synthesis[10] which explains the high number of *SMT_CHM* (80 k) and *SMT_PRO*
74 (67 k) relations.

75 Together, the acquired data allows the extraction of subgraphs corresponding to wildcard triples such
76 as $\langle \text{TiO}_2, \text{CHM_PRO}, ? \rangle$, which correspond to the customized query: “what are the properties of
77 TiO_2 ?”. Further, by accounting for the co-occurrence frequency, a confidence score can be assigned
78 to each triple as is visually represented in Fig 1(a, b) where the applications and phase labels of
79 TiO_2 are separately extracted and presented as individual bipartite graphs such that the size of the
80 node is proportional to $v(\text{TiO}_2; e)$. We see that the most common symmetry/phase labels associated
81 with TiO_2 are ‘rutile’ and ‘anastase’, while the most frequent applications are as electrodes, catalyts
82 and for coating. These are in agreement with the widely available literature on the material[24].
83 There is much less information on CdTe by comparison (18153 vs 1500 triples), but Fig 1(c, d)
84 extracted from MatKG still enables a high-level understanding with some specificity, such as the
85 knowledge that CdTe is used in solar cells and electrodes, and is an optical material as deduced from
86 its properties[25].

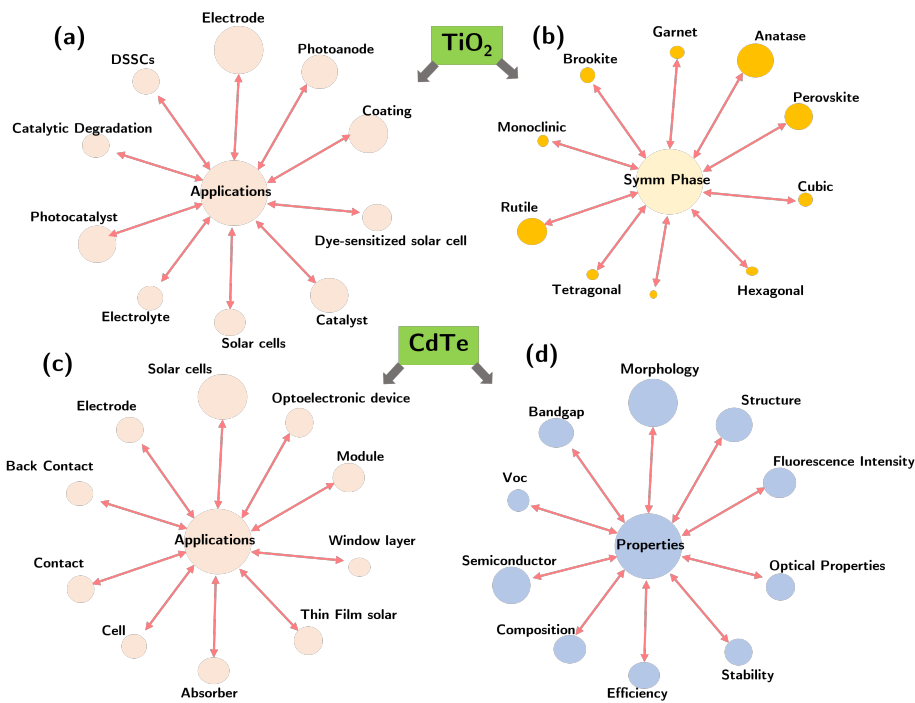


Figure 1: (a) Applications and (b) Symmetry Phase Labels of TiO_2 . (c) Applications and (d) Properties of CdTe . The size of the node is proportional to the co-occurrence frequency of the link.

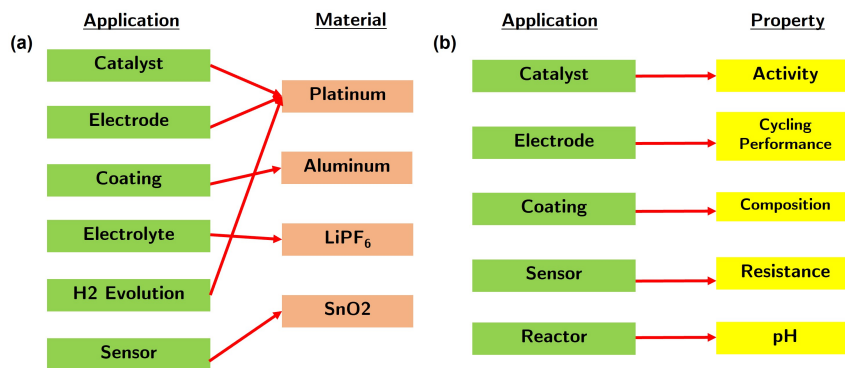


Figure 2: Partitioned (a) application-material (b) application-property subgraphs showing the highest weighted material and property for some select applications

87 In addition to material specific queries, MatKG can be partitioned into relation specific subgraphs such
 88 as the application-material and application – property graphs in Fig 2 (a-b), which shows the highest
 89 weighted material and property respectively for some select applications. Platinum, perhaps the most
 90 widely used metallic catalyst, appears with both ‘catalyst’ and ‘hydrogen evolution’. Aluminum
 91 is identified as a coating material while LiPF_6 is seen to be an electrolyte, both of which are well
 92 known applications of each respectively. In Fig 2 (b) the most common property associated with
 93 electrodes is ‘cycling performance’, while that of catalyst is ‘activity’. Both are in accordance with
 94 our understanding of these concepts. Therefore, MatKG allows the curated visualization and querying
 95 of materials specific data directly extracted from literature at unprecedented scale and resolution.

Table 1: Selected entities and their similarities, demonstrating semantic convergence at the embedding level

Entities	Similarity
(qspr, quantitative structure property relationship)	0.90
(qmom, quadrature method of moments)	0.91
(electromagnetic acoustic resonance, emar)	0.89
(ner, net energy ratio)	0.92
(let, linear energy transfer)	0.91

96 3.2 Embedding representation learning

97 The TransE[19] model with 150 dimensions is found to have the highest MRR (0.49) on the test set.
98 This model was chosen for discovering new links and for performing entity disambiguation.

99 **Entity Linking:** The similarity between embeddings can be used as a measure of the semantic
100 similarity between entities, in turn becoming a useful tool for both co-reference resolution as well
101 as similar – chemical mapping. As shown in Table 1, several pairs of entities such as ‘qspr’ and
102 ‘quantitative structure property relationship’, or ‘ner’ and ‘net energy ratio’ occupy almost identical
103 positions in MatKG and consequently have very similar graph embeddings. This suggests that they
104 are the same semantic token, even though their lexical distance can be substantial. This form of
105 co-reference resolution is currently not an easy task, especially for the sciences[26].

106 **Link Prediction:** Finally, the KGE model was used to make new link predictions between existing
107 entities in the graph. In this way, the model can be used to discover new applications and properties of
108 existing materials, new properties that can be useful to a given application, or a new characterization
109 method for an existing property, etc. This results in a fuller and more integrated knowledge graph,
110 allowing a holistic analysis of structure-property-processing relations, even when such data is absent
111 in the training literature.

112 While the MRR and hits@(1,10,100) are good measures of link predictiveness of the model[27], it
113 is desirable to quantify this inference ability further. To this end, 150 random link predictions were
114 generated by the model across all relationship categories. The top three entities with the highest score
115 for each prediction is manually ranked according to the following criteria: Rank1 if the relationship
116 can be classified as of type *SKOS: Narrow*, Rank 2 if it is of type *SKOS: Broad*, and 3 otherwise,
117 where ‘narrow’ and ‘broad’ are ontological schema specified in SKOS [8]. An example triple is
118 shown in Table 2, Appendix, which lists the top three model predictions for the applications of Fe₂O₃.
119 Some lithium-ion batteries use lithium-iron-oxide as an electrode, which is usually made by the
120 solid-state reaction of Li₂CO₃ and Fe₂O₃, which could explain the first prediction. Since ‘lithium-ion
121 batteries’ is not a direct application of Fe₂O₃, this triple is ranked 2. However, ‘air batteries’ directly
122 use iron/iron oxide as an electrode[28] and hence this triple is assigned rank 1.

123 Of the 150 x 3 predictions made by the model, 47 % were found to have a rank 1, 29 % had a rank of 2,
124 and the rest had a rank of 3 (See Appendix, Table 3 for examples). The utility of this approach is seen
125 in Fig 3, Appendix where previously empty application and characterization subgraphs of Bismuth
126 Telluride (as extracted from MatKG) are populated with meaningful entities through successful link
127 prediction.

128 4 Broader Impact

129 MatKG is the first step towards the complete synthesis of materials knowledge that allows for the
130 richer databases not just for materials but also for applications, properties, and characterization
131 methods. The ability to predict new links between entities in the graph allows the discovery of new
132 materials for existing applications and properties, in finding new applications of existing materials,
133 and novel correlations between synthesis, characterizations and properties. Consequently, MatKG
134 has broad impact for all the three categories of **AI-guided materials design**, **Automated Synthesis**
135 and for **Automated Characterization** .

References

- 136
- 137 [1] D William and J Callister. The structure of crystalline solids. In *Materials Science and*
138 *Engineering*, pages 64–65. Wiley, 1989.
- 139 [2] Vahe Tshitoyan, John Dagdelen, Leigh Weston, Alexander Dunn, Ziqin Rong, Olga Kononova,
140 Kristin A Persson, Gerbrand Ceder, and Anubhav Jain. Unsupervised word embeddings capture
141 latent knowledge from materials science literature. *Nature*, 571(7763):95–98, 2019.
- 142 [3] Alden Dima, Sunil Bhaskarla, Chandler Becker, Mary Brady, Carelyn Campbell, Philippe
143 Dessauw, Robert Hanisch, Ursula Kattner, Kenneth Kroenlein, Marcus Newrock, et al. Infor-
144 matics infrastructure for the materials genome initiative. *Jom*, 68(8):2053–2064, 2016.
- 145 [4] Juan J de Pablo, Nicholas E Jackson, Michael A Webb, Long-Qing Chen, Joel E Moore, Dane
146 Morgan, Ryan Jacobs, Tresa Pollock, Darrell G Schlom, Eric S Toberer, et al. New frontiers for
147 the materials genome initiative. *npj Computational Materials*, 5(1):1–23, 2019.
- 148 [5] Juan J De Pablo, Barbara Jones, Cora Lind Kovacs, Vidvuds Ozolins, and Arthur P Ramirez.
149 The materials genome initiative, the interplay of experiment, theory and computation. *Current*
150 *Opinion in Solid State and Materials Science*, 18(2):99–117, 2014.
- 151 [6] James P McCusker, Neha Keshan, Sabbir Rashid, Michael Deagen, Cate Brinson, and Debo-
152 rah L McGuinness. Nanomine: A knowledge graph for nanocomposite materials science. In
153 *International Semantic Web Conference*, pages 144–159. Springer, 2020.
- 154 [7] Loïc M Roch, Florian Häse, Christoph Kreisbeck, Teresa Tamayo-Mendoza, Lars PE Yunker,
155 Jason E Hein, and Alán Aspuru-Guzik. Chemos: orchestrating autonomous experimentation.
156 *Science Robotics*, 3(19):eaat5559, 2018.
- 157 [8] Alistair Miles and Sean Bechhofer. Skos simple knowledge organization system reference.
158 *W3C recommendation*, 2009.
- 159 [9] David Nadeau and Satoshi Sekine. A survey of named entity recognition and classification.
160 *Linguisticae Investigationes*, 30(1):3–26, 2007.
- 161 [10] Edward Kim, Kevin Huang, Stefanie Jegelka, and Elsa Olivetti. Virtual screening of inorganic
162 materials synthesis parameters with deep learning. *npj Computational Materials*, 3(1):1–9,
163 2017.
- 164 [11] Vineeth Venugopal, Sourav Sahoo, Mohd Zaki, Manish Agarwal, Nitya Nand Gosvami, and
165 NM Anoop Krishnan. Looking through glass: Knowledge discovery from materials science
166 literature using natural language processing. *Patterns*, 2(7):100290, 2021.
- 167 [12] Vineeth Venugopal, Scott R Broderick, and Krishna Rajan. A picture is worth a thousand words:
168 applying natural language processing tools for creating a quantum materials database map. *MRS*
169 *Communications*, 9(4):1134–1141, 2019.
- 170 [13] Vineeth Venugopal, Suresh Bishnoi, Sourabh Singh, Mohd Zaki, Hargun Singh Grover, Mathieu
171 Bauchy, Manish Agarwal, and NM Anoop Krishnan. Artificial intelligence and machine learning
172 in glass science and technology: 21 challenges for the 21st century. *International journal of*
173 *applied glass science*, 12(3):277–292, 2021.
- 174 [14] Leigh Weston, Vahe Tshitoyan, John Dagdelen, Olga Kononova, Amalie Trewartha, Kristin A
175 Persson, Gerbrand Ceder, and Anubhav Jain. Named entity recognition and normalization
176 applied to large-scale information extraction from the materials science literature. *Journal of*
177 *chemical information and modeling*, 59(9):3692–3702, 2019.
- 178 [15] Valentin I Spitzkovsky and Angel X Chang. A cross-lingual dictionary for english wikipedia
179 concepts. 2012.
- 180 [16] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards,
181 Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, et al. Commentary:
182 The materials project: A materials genome approach to accelerating materials innovation. *APL*
183 *materials*, 1(1):011002, 2013.
- 184 [17] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary
185 Ives. Dbpedia: A nucleus for a web of open data. In *The semantic web*, pages 722–735. Springer,
186 2007.

- 187 [18] Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. Yago: a core of semantic knowl-
 188 edge. In *Proceedings of the 16th international conference on World Wide Web*, pages 697–706,
 189 2007.
- 190 [19] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko.
 191 Translating embeddings for modeling multi-relational data. In *NIPS*, pages 2787–2795, 2013.
- 192 [20] Bishan Yang, Scott Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities
 193 and relations for learning and inference in knowledge bases. In *ICLR*, 2015.
- 194 [21] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard.
 195 Complex embeddings for simple link prediction. In *ICML*, pages 2071–2080, 2016.
- 196 [22] Hongyun Cai, Vincent W Zheng, and Kevin Chen-Chuan Chang. A comprehensive survey of
 197 graph embedding: Problems, techniques, and applications. *IEEE Transactions on Knowledge
 198 and Data Engineering*, 30(9):1616–1637, 2018.
- 199 [23] Luca Costabello, Sumit Pai, Chan Le Van, Rory McGrath, Nicholas McCarthy, and Pedro
 200 Tabacof. AmpliGraph: a Library for Representation Learning on Knowledge Graphs, 2019.
 201 URL <https://doi.org/10.5281/zenodo.2595043>.
- 202 [24] Qing Guo, Chuanyao Zhou, Zhibo Ma, and Xueming Yang. Fundamentals of tio2 photocatalysis:
 203 concepts, mechanisms, and challenges. *Advanced Materials*, 31(50):1901997, 2019.
- 204 [25] SH Shin, J Bajaj, LA Moudy, and DT Cheung. Characterization of te precipitates in cdte
 205 crystals. *Applied Physics Letters*, 43(1):68–70, 1983.
- 206 [26] Ozlem Uzuner, Andreea Bodnari, Shuying Shen, Tyler Forbush, John Pestian, and Brett R
 207 South. Evaluating the state of the art in coreference resolution for electronic medical records.
 208 *Journal of the American Medical Informatics Association*, 19(5):786–791, 2012.
- 209 [27] Vivek Khetan, Erin Wetherley, Elena Eneva, Shubhashis Sengupta, Andrew E Fano, et al.
 210 Knowledge graph anchored information-extraction for domain-specific insights. *arXiv preprint
 211 arXiv:2104.08936*, 2021.
- 212 [28] J Requies, MB Güemez, S Perez Gil, VL Barrio, JF Cambra, U Izquierdo, and PL Arias. Natural
 213 and synthetic iron oxides for hydrogen storage and purification. *Journal of Materials Science*,
 214 48(14):4813–4822, 2013.

215 Checklist

216 For all authors...

- 217 1. Do the main claims made in the abstract and introduction accurately reflect the paper’s
 218 contributions and scope? [Yes]
- 219 2. Did you describe the limitations of your work? [Yes]
- 220 3. Did you discuss any potential negative societal impacts of your work? [Yes]
- 221 4. Have you read the ethics review guidelines and ensured that your paper conforms to them?
 222 [Yes]

223 If you are including theoretical results...

- 224 1. Did you state the full set of assumptions of all theoretical results? [N/A]
- 225 2. Did you include complete proofs of all theoretical results? [N/A]

226 If you ran experiments...

- 227 1. Did you include the code, data, and instructions needed to reproduce the main experimental
 228 results (either in the supplemental material or as a URL)? [No] The data is open-source and
 229 freely available. The code and project are still a work-in-progress and code will be released
 230 upon full publication of this work at a later date.
- 231 2. Did you specify all the training details (e.g., data splits, hyperparameters, how they were
 232 chosen)? [No] This will done upon the full publication of this work.

- 233 3. Did you report error bars (e.g., with respect to the random seed after running experiments
234 multiple times)? [N/A]
- 235 4. Did you include the total amount of compute and the type of resources used (e.g., type of
236 GPUs, internal cluster, or cloud provider)? [N/A]
- 237 If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
- 238 1. If your work uses existing assets, did you cite the creators? [Yes]
- 239 2. Did you mention the license of the assets? [Yes]
- 240 3. Did you include any new assets either in the supplemental material or as a URL? [N/A]
- 241 4. Did you discuss whether and how consent was obtained from people whose data you're
242 using/curating? [N/A]
- 243 5. Did you discuss whether the data you are using/curating contains personally identifiable
244 information or offensive content? [N/A]
- 245 If you used crowdsourcing or conducted research with human subjects...
- 246 1. Did you include the full text of instructions given to participants and screenshots, if applica-
247 ble? [N/A]
- 248 2. Did you describe any potential participant risks, with links to Institutional Review Board
249 (IRB) approvals, if applicable? [N/A]
- 250 3. Did you include the estimated hourly wage paid to participants and the total amount spent
251 on participant compensation? [N/A]

252 5 Appendix

253 5.1 KG construction

254 For every entity e , the lexical frequency $L(e)$ is defined as the fraction of documents where e is
255 present at least once, where a document could either be one of N_a abstracts or N_c figure captions. For
256 every pair of entities $(e_1; e_2)$ in a given document, a co-occurrence function $CO(e_1; e_2)$ is defined
257 such that:

$$CO(e_1; e_2) = \begin{cases} 1 & \text{if both } e_1 \text{ and } e_2 \text{ present in the document} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

258 The co-occurrence frequency $v(e_1; e_2)$ is then defined as :

$$v(e_1; e_2) = \frac{\sum^{N_a + N_c} CO(e_1; e_2)}{N_a + N_c} \quad (2)$$

259 $v(e_1; e_2)$ therefore is a measure of how many times the given pair of entities $(e_1; e_2)$ co-occur
260 in the document corpus. Subsequently, two approaches are employed to assign a link to $(e_1; e_2)$.
261 Approach (1) is based on the premise that if $v(e_1; e_2) > \frac{L(e_1) * L(e_2)}{(N_a + N_c)^2}$, then the entities e_1 and e_2 are
262 strongly correlated as they occur far more often than their conditional probabilities allow. Approach
263 (2) however, retains all entity pairs but appends their co-occurrence frequency as a weight in the
264 knowledge representation model

$$Relation(e_1; e_2) = \begin{cases} \langle e_1; T[e_1]_T[e_2], e_2 \rangle & v(e_1; e_2) > \beta, \text{ where } \beta \text{ is a threshold} \\ \langle e_1; T[e_1]_T[e_2]; e_2, v(e_1; e_2) \rangle & v(e_1; e_2) > \beta \text{ where } \beta \approx 10 \end{cases}$$

265

Table 2: NER Categories in MatKG and the number of unique entities in each category.

NER Category	Number of Entities
Property (PRO)	27048
Chemical (CHM)	23438
Characterization Method	10908
Synthesis Method	8547
Application	7009

Table 3: Selected relationships and their instance count in MatKG.

Relationship	Number of Triple
<i>CHM_CHM</i>	499994
<i>PRO_PRO</i>	368381
<i>CHM_PRO</i>	252714
<i>PRO_DSC</i>	146929
<i>CMT_CHM</i>	141955
<i>CHM_DSC</i>	139740
<i>CMT_PRO</i>	108233
<i>CMT_CMT</i>	100675
<i>APL_PRO</i>	91466
<i>CHM_APL</i>	89117
<i>CHM_SMT</i>	80349

Table 4: Model predictions for the triple $\langle \text{Fe}_2\text{O}_3, \text{CHM_PRO}, X \rangle$ where X is a property. The triples are ranked according to the scheme described in Results

Subject	relationship	Object	Rank
Fe_2O_3	<i>CHM_APL</i>	<i>lithium ion batteries</i>	2
Fe_2O_3	<i>CHM_APL</i>	<i>electrocatalysts</i>	1
Fe_2O_3	<i>CHM_APL</i>	<i>air batteries</i>	1

Table 5: Top three model predicted links for selected examples with model score, custom rank, and cited doi

Subject	relationship	Object	Score	Rank	Citation url
optical material	<i>APL_CHM</i>	In_2O_3	5.5	1	https://en.wikipedia.org/wiki/Indium(III)_oxide
optical material	<i>APL_CHM</i>	CdO	5.27	1	https://en.wikipedia.org/wiki/Cadmium_oxide
optical material	<i>APL_CHM</i>	Zinc Oxide	5.26	1	https://en.wikipedia.org/wiki/Zinc_oxide
anodic electrode	<i>APL_CHM</i>	Graphite	3.00	1	10.1016/j.ensm.2020.12.027
anodic electrode	<i>APL_CHM</i>	Carbon-fiber	3.00	1	10.1016/C2015-0-00574-3
anodic electrode	<i>APL_CHM</i>	LiClO_4	2.90	2	https://en.wikipedia.org/wiki/Lithium_perchlorate
nuclear reactor	<i>APL_CHM</i>	Beryllium	7.02	1	https://www.energy.gov/ehss/about-beryllium
nuclear reactor	<i>APL_CHM</i>	Carbide	6.41	2	https://en.wikipedia.org/wiki/Uranium_carbide
nuclear reactor	<i>APL_CHM</i>	Tungsten	6.38	1	10.1016/j.ijhydene.2016.02.019
smes	<i>APL_PRO</i>	dmain	0.34	3	N/A
smes	<i>APL_PRO</i>	transmitted current	0.28	1	https://en.wikipedia.org/wiki/Superconducting_magnetic_energy_storage
smes	<i>APL_PRO</i>	u11	0.20	3	N/A
reverse water gas shift reaction	<i>APL_CHM</i>	$\text{C}_6\text{H}_5\text{OH}$	5.22	3	N/A
reverse water gas shift reaction	<i>APL_CHM</i>	Naphtha	5.22	3	N/A
reverse water gas shift reaction	<i>APL_CHM</i>	diethylether	4.79	3	N/A

