

# Commonsense Ontology Micropatterns

Andrew Eells,<sup>1</sup> Brandon Dave,<sup>2</sup> Pascal Hitzler,<sup>1</sup> Cogan Shimizu<sup>2</sup>

<sup>1</sup> Kansas State University, Manhattan, KS, USA

<sup>2</sup> Wright State University, Dayton, OH, USA

**Abstract.** The previously introduced Modular Ontology Modeling methodology (MOMo) attempts to mimic the human analogical process by using modular patterns to assemble more complex concepts. To support this, MOMo organizes ontology design patterns into design libraries, which are programmatically queryable, to support accelerated ontology development, for both human and automated processes. However, a major bottleneck to large-scale deployment of MOMo is the (to-date) limited availability of ready-to-use ontology design patterns. At the same time, Large Language Models have quickly become a source of common knowledge and, in some cases, replacing search engines for questions. In this paper, we thus present a collection of 104 ontology design patterns representing often occurring nouns, curated from the common-sense knowledge available in LLMs, organized into a fully-annotated modular ontology design library ready for use with MOMo.

## 1 Introduction

Humans perform transfer learning through analogous connections in understanding the unknowns of the world with previously acquired knowledge [6,7]. In the scope of problem-solving, humans find similarities in their current problem to draw logical connections that may aid in finding a working solution, which, regardless of success, becomes part of the learning process for the next iteration of problem-solving. The task of drawing similarities requires humans to utilize their logical skills for generalizing a problem or concept to allow for adaptations from their acquired knowledge.

Neurosymbolic Artificial Intelligence (AI) systems, in their pursuit of General AI, also seek this ability to conduct transfer and analogy learning. We posit that the use of modular ontology is one way that this can be accomplished [11]. In particular, the notion of starting with a simple and (perhaps) well-understood pattern of knowledge and adapting it to new scenarios and use-cases; the mechanism by which we accomplish this we call *modularization*. This overarching methodology is known as Modular Ontology Modeling (MOMo), and it is by now an established paradigm for constructing knowledge graphs (KG) and their schemas (i.e., an ontology) [14].

Ontology design patterns, which are tiny ontologies that are intended to solve domain-invariant modeling problems [5], are the usual starting point for the MOMo methodology. They may be sourced from many places, including

the Ontology Design Pattern Portal<sup>3</sup> or Modular Ontology Design Libraries (MODLs) [15]. MODLs are annotated, curated collections of ODPs, combined with an index ontology that can be queried for patterns of different types and relationships between patterns.

In this paper, we provide such a design library, which we call CS-MODL, composed of *commonsense micropatterns*. For the purposes of this work, we define our title.

- **commonsense** we loosely mean “what first comes to mind.” In this case, we ask an LLM about common properties of common nouns (as discussed below) and receive some simple and naïve responses, which are then encoded into:
- **micropatterns** we contrast micropatterns versus ODPs. In this case, micropatterns do not have a sophisticated or rich semantics. Indeed, they are defined using `rdfs:domain` and `rdfs:range`, without more complicated OWL axioms.

Concretely, we describe the methodology we used to construct the 104 commonsense micropatterns contained in CS-MODL, discuss an example, and provide the library as an online resource.

The rest of the paper is organized as follows: Section 2 introduces our methodologies for obtaining the appropriate information from our chosen LLM and constructing the commonsense micropatterns. Section 3 provides an example of a generated commonsense micropattern. Section 4 summarizes the contributions of our commonsense library of patterns. Section 5 wraps up the paper with notes on future areas we observe as next steps to this research.

## 2 Methodology

The construction of these commonsense patterns was completed in two steps. First, we prompt an LLM in several ways (Section 2.3) to produce an ontology design pattern for each noun in question (Section 2.2). Then, the responses are consolidated into patterns via a set of heuristics (Section 2.4), and the patterns are packaged into a MODL (Section 4).

### 2.1 Large Language Models as Commonsense

LLMs, like GPT-4, ingest and learn from more information than any human possibly could. While it is certainly arguable that any LLM demonstrates human-level intelligence, it is also inarguable that it can present a commonsense view of the world – in particular through the a declarative memory, having learned, perhaps as an accident of learning from “mostly correct” language, facts about the world [2].

They are able to represent this information and “knowledge” in human-readable ways with surprising fluency, and as we demonstrate below: frequently

<sup>3</sup> <https://ontologydesignpatterns.org/>

1. Air	27. Force	53. Man	79. Research
2. Area	28. Friend	54. Member	80. Result
3. Art	29. Game	55. Minute	81. Right
4. Back	30. Girl	56. Moment	82. Room
5. Body	31. Government	57. Money	83. School
6. Book	32. Group	58. Month	84. Service
7. Business	33. Guy	59. Morning	85. Side
8. Car	34. Hand	60. Mother	86. Sofa
9. Case	35. Head	61. Name	87. State
10. Chair	36. Health	62. Night	88. Story
11. Change	37. History	63. Number	89. Student
12. Child	38. Home	64. Office	90. Study
13. City	39. Hour	65. Others	91. System
14. Community	40. House	66. Parent	92. Teacher
15. Company	41. Idea	67. Part	93. Team
16. Couch	42. Information	68. Party	94. Thing
17. Country	43. Issue	69. People	95. Time
18. Day	44. Job	70. Person	96. War
19. Door	45. Kid	71. Place	97. Water
20. Education	46. Kind	72. Point	98. Way
21. End	47. Law	73. Power	99. Week
22. Eye	48. Level	74. President	100. Woman
23. Face	49. Life	75. Problem	101. Word
24. Fact	50. Line	76. Program	102. Work
25. Family	51. Lot	77. Question	103. World
26. Father	52. Loveseat	78. Reason	104. Year

**Fig. 1.** The 104 nouns conceptualized into our commonsense micropatterns.

in valid RDF [3]. The following sections describe how we decided on which facts to extract (i.e., which commonsense notions to prompt), how we prompted for them, and the resultant collection of patterns.

## 2.2 Noun Selection

The nouns selected are the 101 most common nouns in American English, as listed in the Corpus of Contemporary American English (COCA) [1]. We additionally include the nouns “chair”, “sofa”, and “loveseat” from our preliminary explorations for viability of our idea. All 104 nouns in our investigation are listed in Figure 1.

### 2.3 Prompt Engineering

We prompted the LLM<sup>4</sup> in several different ways, as we wished to test how it would respond to multiple variants of the same prompts. The initial prompts were a command (1) and a request (2), and then variations were added to include the noun as “the following” (3 and 4), as seen below.

1. Generate an ontology that covers [noun].
2. Could you develop a basic design pattern for representing [noun] in an ontology?
3. Generate an ontology that covers the following: [noun].
4. Could you develop a basic design pattern for representing the following in an ontology: [noun].

Finally, the LLM was prompted to generate 9 additional variations of each prompt. See below for Prompt 1’s variations. In each case, the prompt was tried with and without the additional instruction of “Provide it in valid Turtle/RDF format, excluding any extra text.” for a total of 20 prompts per noun.

1. Generate an ontology that covers [noun].
2. Develop an ontology dedicated to [noun].
3. Construct an ontology focused on [noun].
4. Build an ontology surrounding the concept of a [noun].
5. Formulate an ontology related to [noun].
6. Establish an ontology based on [noun].
7. Design an ontology to encompass [noun].
8. Produce an ontology specifically for [noun].
9. Compose an ontology to represent [noun].
10. Make an ontology that pertains to [noun].

### 2.4 Response Consolidation

As noted, the LLM generated 80 responses for each of the 104 nouns in our collection. These responses were stored in a tab-delimited (TSV) file. However, not every response contained valid RDF and some contained no RDF at all. The number of responses was too large for manual processing, so we developed a script that would apply heuristics to each response file in order to integrate the responses into a single micropattern. The script consists of three parts, applied to each noun:

1. **extraction** – for each specific prompt response (i.e., a row in the TSV), the script attempts to identify and extract the portion of the response in RDF. Unfortunately, not every response output an ODP or easily recognizable RDF. Generally speaking, cleaning required the removal of extra text (even when prompted *not* to provide extra text) and surrounding Markdown formatting (e.g., code fencing).

---

<sup>4</sup> For this work, we used GPT-4-0613.

2. **cleaning** – a set of heuristics was applied to each extract RDF (generally in turtle format) to allow for ingestion of the RDF into an `rdflib`<sup>5</sup> graph structure. The most common type of cleaning necessary was the usage of newline characters. The LLM understood how to comment (i.e., separating class declarations from properties from assertions) but did not understand new lines. This frequently resulted in malformed Turtle, since the first part of the triple would be commented out.
3. **integration** – once each valid (or cleanable) TTL file was extracted from the response, we combine the properties (removing all triples pertaining to individuals), and annotate the result, giving us a commonsense micropattern for the noun. The annotations are further discussed in Section 4.

While we note that we combine all (valid and cleanable) responses into a single micropattern, the script is set up in a much more flexible way.

First, it is necessary to recall that the process by which ODPs are utilized in the MOMo process is based on customization to the use-case at hand; that is properties are frequently removed, added, or renamed for best fit. This process is known as *template-based instantiation* [8], and it is a core aspect of the methodology. Specifically, each property with (explicit or inferred) domain and range, is assessed for fit. If it is not applicable (i.e., out of scope or too specific or general), the property and associated sub-graphs are removed. Conversely, if there are parts of the use-case that are not covered by the ODP, these must be added. In automated processes, it is generally easier to pare down an overly superfluous pattern. That is, if there are no matches or fillers for the slots that a pattern provides, they should be dropped. We follow this strategy for our micropatterns, meaning that we include each property from every response, which includes variants of the same property. For example, in Figure 2, which shows a schema diagram of a commonsense micropattern for `Air`, `hasHumidity` has a range of both `Humidity` and `xsd:float`. It is conceivable that one or the other is unnecessary, or that they may be used in tandem (i.e., where a connection from `Humidity` is made to the `xsd:string`), e.g., as a type of shortcut or view [13,17].

Yet, this strategy may not be applicable for all systems. As such, the script allows for a voting mechanism to take place. This might be useful if additional prompt variants are used or if multiple LLMs are prompted for ODPs, thus allowing for properties (with domains and ranges) that appear in multiple patterns from multiple sources are those that appear in the final micropattern. The voting mechanism is as follows.

For the usable response, we extracted 3-tuples expressing Properties associated with a noun, and the Property’s respective described scoped domain and range. This process also included the notion of a voting process, which would decide the final noun pattern’s implemented properties based on occurrence per generated noun ontology. Our noun ontologies ended up with limited duplicate occurrences of the same triplets (s, p, o) (ranging from 1 to 4 total occurrences), and so, we chose not to implement the voting process and a threshold to implement properties to the final noun pattern. This allows us to maximally express

<sup>5</sup> <https://rdflib.readthedocs.io/en/stable/>



### 3 Example Pattern

We have selected a pattern, `Air`, which exemplifies many interesting properties of our commonsense micropatterns. The schema diagram for this micropattern is shown in Figure 2. The gold boxes reflect classes. The yellow ellipses indicate datatypes. Black filled arrows are binary relations. When labels overlap multiple arrows, we use this to show that there are multiple ranges for the property.

As noted in Section 2.4, there are some reuse of property names (e.g., `hasHumidity`) that would technically be not allowed in OWL. However, we suggest that *pruning* is easier than growing, so a system that utilizes CS-MODL can simply remove the unused property, as necessary.

The turtle file for this example can be found in the repository.

#### 3.1 License & Availability

Our artifacts are made available online through a GitHub repository.<sup>6</sup> All data (which includes documentation, prompts, responses, extracted and cleaned TTL files, integrated patterns, and CS-MODL itself) and the supporting scripts that generated the data are provided under an open and permissive license (GNU Lesser General Public License v2.1), and it is also included explicitly in the repository.

## 4 A Commonsense Modular Ontology Design Library

The purpose of a MODL is to collect together (curate) a set of ontology design patterns. In the first version [15], these were generally well-documented and useful patterns. A forthcoming second version has many more, spanning a wider set of application scenarios.<sup>7</sup> This particular MODL, CS-MODL, is similar insofar that it also satisfies broad use-cases. The purpose of which is to collect together patterns of similar quality, abstraction, and expressivity for the purpose of modeling commonly occurring entities (e.g., in informal language or the real world).

The CS-MODL, like other MODLs, is programmatically queryable through OPaL annotations [10,9] indicating that CS-MODL contains sufficient metadata scaffolding such that top-level conceptual components (i.e., the key notion) of patterns can be easily matched to external entities detected by a framework.<sup>8</sup> For example, if a vision interpreter finds something with four legs, it could match to the `Chair` pattern. Inversely, if a chair is detected, *most* chairs have four legs, and this can be inferred from the matched pattern. This matching process can also be used for inferring *behaviors* or *functionalities* of said items, which is even more applicable to this use case.

<sup>6</sup> <https://github.com/kastle-lab/commonsense-micropatterns/>

<sup>7</sup> <https://github.com/kastle-lab/modular-ontology-design-library>

<sup>8</sup> We note that we are trying to have a better acronym – OPaL – instead of the previously published acronym!

## 5 Conclusion

Large language models have quickly become a source of commonsense information, in some cases supplanting search engine use. This easily available source of knowledge can enhance transfer and analogy learning, especially in combination with knowledge engineering methodologies (such as modular ontology modeling).

In this paper, we have described our methodology for obtaining 104 commonsense micropatterns from an LLM. This process is generalizable from our chosen list, and the code to produce additional patterns is available online and permissibly licensed. These patterns are organized into a modular ontology design library (MODL), which we call CS-MODL, also available in the same repository. CS-MODL can facilitate the construction of internal knowledge frameworks of automated systems through specializing or generalizing from the commonsense micropatterns.

### Future Work

We have identified some areas for future enhancement.

1. The list of nouns that we utilized to seed our prompts for patterns is naïve. That is, no effort was put into organizing them into a subsumption hierarchy or otherwise identifying unifying concepts. Creating both this hierarchy, as well as examining similar hierarchies for relations between these concepts, is an immediate next step. Additionally, these relations can be encoded in the CS-MODL index allowing for substantially more sophisticated pattern identification.
2. We will evaluate the use of CS-MODL in a more principled way (i.e., through use in a case study) pertaining to ontology learning from text, and in a multi-agent teaming, visual sensor fusion exercise.
3. Describe the patterns more expressively (i.e., fully describe the patterns in OWL instead of only using RDFS).
4. Improved documentation: schema diagrams, summaries, and explanation of the formalization (i.e., in the style of [16] and [12]) will be explored.

*Acknowledgement.* The authors acknowledge partial funding by the National Science Foundation under grant 2333532 "Proto-OKN Theme 3: An Education Gateway for the Proto-OKN."



## References

1. English-Corpora: COCA, <https://www.english-corpora.org/coca/>
2. Cohen, R., Geva, M., Berant, J., Globerson, A.: Crawling the internal knowledge-base of language models (2023)
3. Cyganiak, R., Wood, D., Lanthaler, M. (eds.): RDF 1.1 Concepts and Abstract Syntax. W3C Recommendation 25 February 2014 (2014), available from <http://www.w3.org/TR/rdf11-concepts/>
4. Eberhart, A., Shimizu, C., Chowdhury, S., Sarker, M.K., Hitzler, P.: Expressibility of OWL axioms with patterns. In: Verborgh, R., Hose, K., Paulheim, H., Champin, P., Maleshkova, M., Corcho, Ó., Ristoski, P., Alam, M. (eds.) The Semantic Web - 18th International Conference, ESWC 2021, Virtual Event, June 6-10, 2021, Proceedings. Lecture Notes in Computer Science, vol. 12731, pp. 230–245. Springer (2021). [https://doi.org/10.1007/978-3-030-77385-4\\_14](https://doi.org/10.1007/978-3-030-77385-4_14), [https://doi.org/10.1007/978-3-030-77385-4\\_14](https://doi.org/10.1007/978-3-030-77385-4_14)
5. Gangemi, A., Presutti, V.: Ontology design patterns. In: Staab, S., Studer, R. (eds.) Handbook on Ontologies, pp. 221–243. International Handbooks on Information Systems, Springer (2009). [https://doi.org/10.1007/978-3-540-92673-3\\_10](https://doi.org/10.1007/978-3-540-92673-3_10), [https://doi.org/10.1007/978-3-540-92673-3\\_10](https://doi.org/10.1007/978-3-540-92673-3_10)
6. Gentner, D., Smith, L.: Analogical reasoning. In: Ramachandran, V. (ed.) Encyclopedia of Human Behavior (Second Edition), pp. 130–136. Academic Press, San Diego, second edition edn. (2012). <https://doi.org/https://doi.org/10.1016/B978-0-12-375000-6.00022-7>, <https://www.sciencedirect.com/science/article/pii/B9780123750006000227>
7. Gentner, D., Smith, L.A.: Analogical learning and reasoning., pp. 668–681. Oxford library of psychology., Oxford University Press, New York, NY, US (2013). <https://doi.org/10.1093/oxfordhb/9780195376746.001.0001>, <https://doi.org/10.1093/oxfordhb/9780195376746.001.0001>
8. Hammar, K., Presutti, V.: Template-based content ODP instantiation. In: Hammar, K., Hitzler, P., Krisnadhi, A., Lawrynowicz, A., Nuzzolese, A.G., Solanki, M. (eds.) Advances in Ontology Design and Patterns [revised and extended versions of the papers presented at the 7th edition of the Workshop on Ontology and Semantic Web Patterns, WOP@ISWC 2016, Kobe, Japan, 18th October 2016]. Studies on the Semantic Web, vol. 32, pp. 1–13. IOS Press (2016). <https://doi.org/10.3233/978-1-61499-826-6-1>, <https://doi.org/10.3233/978-1-61499-826-6-1>
9. Hirt, Q., Shimizu, C., Hitzler, P.: Extensions to the ontology design pattern representation language. In: Janowicz, K., Krisnadhi, A.A., Poveda-Villalón, M., Hammar, K., Shimizu, C. (eds.) Proceedings of the 10th Workshop on Ontology Design and Patterns (WOP 2019) co-located with 18th International Semantic Web Conference (ISWC 2019), Auckland, New Zealand, October 27, 2019. CEUR Workshop Proceedings, vol. 2459, pp. 76–75. CEUR-WS.org (2019), <http://ceur-ws.org/Vol-2459/short2.pdf>
10. Hitzler, P., Gangemi, A., Janowicz, K., Krisnadhi, A.A., Presutti, V.: Towards a simple but useful ontology design pattern representation language. In: Blomqvist, E., Corcho, Ó., Horridge, M., Carral, D., Hoekstra, R. (eds.) Proceedings of the 8th Workshop on Ontology Design and Patterns (WOP 2017) co-located with the 16th International Semantic Web Conference (ISWC 2017), Vienna, Austria, October 21, 2017. CEUR Workshop Proceedings, vol. 2043. CEUR-WS.org (2017), <http://ceur-ws.org/Vol-2043/paper-09.pdf>

11. Hitzler, P., Shimizu, C.: Modular ontologies as a bridge between human conceptualization and data. In: Chapman, P., Endres, D., Pernelle, N. (eds.) Graph-Based Representation and Reasoning - 23rd International Conference on Conceptual Structures, ICCS 2018, Edinburgh, UK, June 20-22, 2018, Proceedings. Lecture Notes in Computer Science, vol. 10872, pp. 3-6. Springer (2018). [https://doi.org/10.1007/978-3-319-91379-7\\_1](https://doi.org/10.1007/978-3-319-91379-7_1), [https://doi.org/10.1007/978-3-319-91379-7\\_1](https://doi.org/10.1007/978-3-319-91379-7_1)
12. Karima, N., Hammar, K., Hitzler, P.: How to document ontology design patterns. In: Hammar, K., Hitzler, P., Lawrynowicz, A., Krisnadhi, A., Nuzzolese, A., Solanki, M. (eds.) Advances in Ontology Design and Patterns, Studies on the Semantic Web, vol. 32, pp. 15-28. IOS Press, Amsterdam (2017)
13. Krisnadhi, A.A., Hitzler, P., Janowicz, K.: On the capabilities and limitations of OWL regarding typecasting and ontology design pattern views. In: Tamma, V.A.M., Dragoni, M., Gonçalves, R.S., Lawrynowicz, A. (eds.) Ontology Engineering - 12th International Experiences and Directions Workshop on OWL, OWLED 2015, co-located with ISWC 2015, Bethlehem, PA, USA, October 9-10, 2015, Revised Selected Papers. Lecture Notes in Computer Science, vol. 9557, pp. 105-116. Springer (2015). [https://doi.org/10.1007/978-3-319-33245-1\\_11](https://doi.org/10.1007/978-3-319-33245-1_11), [https://doi.org/10.1007/978-3-319-33245-1\\_11](https://doi.org/10.1007/978-3-319-33245-1_11)
14. Shimizu, C., Hammar, K., Hitzler, P.: Modular ontology modeling. *Semantic Web* **14**(3), 459-489 (2023). <https://doi.org/10.3233/SW-222886>, <https://doi.org/10.3233/SW-222886>
15. Shimizu, C., Hirt, Q., Hitzler, P.: MODL: A modular ontology design library. In: Janowicz, K., Krisnadhi, A.A., Poveda-Villalón, M., Hammar, K., Shimizu, C. (eds.) Proceedings of the 10th Workshop on Ontology Design and Patterns (WOP 2019) co-located with 18th International Semantic Web Conference (ISWC 2019), Auckland, New Zealand, October 27, 2019. CEUR Workshop Proceedings, vol. 2459, pp. 47-58. CEUR-WS.org (2019), <http://ceur-ws.org/Vol-2459/paper4.pdf>
16. Shimizu, C., Hitzler, P., Hirt, Q., Rehberger, D., Estrecha, S.G., Foley, C., Sheill, A.M., Hawthorne, W., Mixer, J., Watrall, E., Carty, R., Tarr, D.: The enslaved ontology: Peoples of the historic slave trade. *J. Web Semant.* **63**, 100567 (2020). <https://doi.org/10.1016/J.WEBSEM.2020.100567>, <https://doi.org/10.1016/j.websem.2020.100567>
17. Shimizu, C., Stephen, S., Christou, A., Currier, K., Mahdavejad, M.S., Norouzi, S.S., Dalal, A., Barua, A., Fisher, C.K., D'Onofrio, A., Thelen, T., Janowicz, K., Rehberger, D., Schildhauer, M., Hitzler, P.: Knowwhereregraph-lite: A perspective of the knowwhereregraph. In: Ortiz-Rodríguez, F., Villazón-Terrazas, B., Tiwari, S., Bobed, C. (eds.) Knowledge Graphs and Semantic Web - 5th Iberoamerican Conference and 4th Indo-American Conference, KGSWC 2023, Zaragoza, Spain, November 13-15, 2023, Proceedings. Lecture Notes in Computer Science, vol. 14382, pp. 199-212. Springer (2023). [https://doi.org/10.1007/978-3-031-47745-4\\_15](https://doi.org/10.1007/978-3-031-47745-4_15), [https://doi.org/10.1007/978-3-031-47745-4\\_15](https://doi.org/10.1007/978-3-031-47745-4_15)