



HAL
open science

Refined Commonsense Knowledge from Large-Scale Web Contents

Tuan-Phong Nguyen, Simon Razniewski, Julien Romero, Gerhard Weikum

► **To cite this version:**

Tuan-Phong Nguyen, Simon Razniewski, Julien Romero, Gerhard Weikum. Refined Commonsense Knowledge from Large-Scale Web Contents. 2022. hal-03537547v1

HAL Id: hal-03537547

<https://hal.science/hal-03537547v1>

Preprint submitted on 20 Jan 2022 (v1), last revised 10 Sep 2022 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Refined Commonsense Knowledge from Large-Scale Web Contents

Tuan-Phong Nguyen, Simon Razniewski, Julien Romero, Gerhard Weikum

Abstract—Commonsense knowledge (CSK) about concepts and their properties is useful for AI applications. Prior works like ConceptNet, COMET and others compiled large CSK collections, but are restricted in their expressiveness to subject-predicate-object (SPO) triples with simple concepts for S and strings for P and O. This paper presents a method, called ASCENT++, to automatically build a large-scale knowledge base (KB) of CSK assertions, with refined expressiveness and both better precision and recall than prior works. ASCENT++ goes beyond SPO triples by capturing composite concepts with subgroups and aspects, and by refining assertions with semantic facets. The latter is important to express the temporal and spatial validity of assertions and further qualifiers. ASCENT++ combines open information extraction with judicious cleaning and ranking by typicality and saliency scores. For high coverage, our method taps into the large-scale crawl C4 with broad web contents. The evaluation with human judgements shows the superior quality of the ASCENT++ KB, and an extrinsic evaluation for QA-support tasks underlines the benefits of ASCENT++. A web interface, data and code can be accessed at <https://www.mpi-inf.mpg.de/ascentpp>.

Index Terms—Commonsense Knowledge, Knowledge Base Construction



1 INTRODUCTION

Motivation. Commonsense knowledge (CSK) is a long-standing goal of AI [1], [2], [3]: equip machines with structured knowledge about everyday concepts and their properties (e.g., elephants are big and eat plants, buses carry passengers and drive on roads) and about typical human behavior and emotions (e.g., children love visiting zoos, children enter buses to go to school). In recent years, research on the automatic acquisition of CSK assertions has significantly been advanced, and several commonsense knowledge bases (CSKBs) of considerable size have been constructed (see, e.g., [4], [5], [6], [7]). Use cases for CSK include particularly language-centric tasks such as question answering and conversational systems (see, e.g., [8], [9], [10]).

Examples: Question-answering systems often need CSK as background knowledge for robust answers. For example, when a child asks “Which zoos have habitats for T-Rex dinosaurs?”, the system should point out that 1) dinosaurs are extinct, and 2) they can be seen in museums, not in zoos. Dialogue systems should not only generate plausible utterances from a language model but should also be situative, understand metaphors and implicit contexts and avoid blunders. For example, when a user says “tigers will soon join the dinosaurs”, the machine should understand that this refers to an endangered species rather than alive tigers invading museums.

This paper aims to advance the automatic acquisition of CSK assertions from online content to bring better expressiveness, higher precision, and broader coverage.

- *Tuan-Phong Nguyen, Simon Razniewski and Gerhard Weikum are with the Max Planck Institute for Informatics, Germany. Julien Romero is with Telecom SudParis, France.
E-mail: {tuanphong, srazniew, weikum}@mpi-inf.mpg.de, julien.romero@telecom-sudparis.eu*

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

State-of-the-art and its limitations. Large KBs like DBpedia, Wikidata, or Yago largely focus on encyclopedic knowledge on individual entities like people or places and are very sparse on general concepts [11]. Notable projects that focus on CSK include ConceptNet [4], WebChild [5], Mosaic TupleKB [6] and Quasimodo [7]. They are all based on SPO triples as knowledge representation and have major shortcomings:

- *Expressiveness for S:* As subjects, prior CSKBs strongly focus on simple concepts expressed by single nouns (e.g., elephant, car, trunk). This misses semantic refinements (e.g., diesel car vs. electric car) that lead to different properties (e.g., polluting vs. green) and is also prone to word-sense disambiguation problems (e.g., elephant trunk vs. car trunk). Even when CSK acquisition considers multi-word phrases, it still lacks the awareness of semantic relations among concepts. Hypernymy lexicons like WordNet or Wiktionary are also very sparse on multi-word concepts. With these limitations, word-sense disambiguation does not work robustly; prior attempts showed mixed results at best (e.g., [5], [6]).
- *Expressiveness for P and O:* Predicates and objects are treated as monolithic strings, such as:
 - o A1: buses, [used for], [transporting people];
 - o A2: buses, [used for], [bringing children to school];
 - o A3: buses, [carry], [passengers];
 - o A4: buses, [drop], [visitors at the zoo on the weekend].

This misses the equivalence of assertions A1 and A3 and cannot capture the semantic relation between A1 and A2, namely, A2 refining A1. Finally, the spatial facets of A2 and A4 are cluttered into unrelated strings, and the temporal facet in A4 is not explicit either. The alternative of restricting P to a small number of pre-specified predicates (e.g., [4], [5]) and O to concise phrases comes at the cost of much lower coverage.

- *Quality of CSK assertions:* Some of the major CSKBs have prioritized precision (i.e., the validity of the assertions) but have fairly limited coverage (e.g., [4], [6]). Others have broader coverage but include many noisy if not implausible assertions (e.g., [5], [7]). Very few have paid attention to the saliency of assertions, i.e., the degree to which statements are common knowledge, as opposed to merely capturing many assertions. However, projects along these lines (e.g., [4], [12]) fall short in coverage.

ASCENT++ aims to overcome these limitations of prior works while retaining their positive characteristics. In particular, we aim to reconcile high precision with wide coverage and saliency. Like [6], [7], we aim to acquire open assertions (as opposed to pre-specified predicates only) but strive for more expressive representations by refining subjects and capturing semantic facets of assertions. Furthermore, we also provide a canonicalized version of the resulting CSKB in the ConceptNet schema, thus enabling direct use in applications relying on that fixed schema.

Approach. We present the ASCENT++ method for acquiring CSK assertions with refined semantics from web contents. ASCENT++ operates in two phases: (i) scalable extraction from a large web corpus, and (ii) aggregation and consolidation. In the first phase, ASCENT++ processes the C4 crawl [13] a collection of 365 million English web pages. ASCENT++ extracts OpenIE-style tuples by carefully designed dependency-parse-based rules, taking into account assertions for subgroups and aspects of target subjects. The extractor uses cues from prepositional phrases and adverbs to detect semantic facets, and uses supervised classification for eight facet types.

In the second phase, on a per-subject basis, ASCENT++ identifies relevant web pages based on embedding similarity to reference Wikipedia articles, this way being able to distinguish homonyms like “star (celebrity)” vs. “star (celestial body)”. Assertions are iteratively grouped and organized using embedding-based similarity. Open-IE-style assertions are canonicalized into the established ConceptNet schema. Finally, the resulting statements are ranked both by salience and typicality scores, by a supervised machine learning model.

We ran ASCENT++ on the C4 crawl for 10,000 salient concepts from ConceptNet as target subjects. To evaluate the intrinsic quality of the resulting CSKB, we obtained human judgements for a large sample. Our CSKB significantly improves over automatically-built state-of-the-art CSK collections, in terms of precision and recall.

In addition, we performed an extrinsic evaluation in which commonsense knowledge was used to support language models in question answering. Using 3 different settings and 6 different CSKBs, ASCENT++ significantly outperformed language models without this CSK background knowledge in 2 of the 3 settings, and was best or second best among all 6 CSKBs in all 3 settings..

Contributions. Key contributions of this work are the following:

- An expressive model for commonsense knowledge with advanced semantics, with subgroups of subjects and faceted assertions as first-class citizens, and scores for typicality and saliency.

- An unsupervised automated method for populating the model with high-quality CSK assertions by extraction from web contents at large scale and various techniques for aggregation and cleaning.
- Constructing and publicly releasing a high-quality CSKB with 2 million assertions for 10,000 important concepts.

A web interface to the ASCENT++ KB, along with downloadable data and code is accessible at <https://www.mpi-inf.mpg.de/ascentpp>.

Prior publication. This paper substantially extends an earlier conference paper [14]. Major extensions are:

- The extraction method is completely re-worked, as the earlier version operated solely on top-ranked query results from search engine APIs whereas ASCENT++ processes a huge web crawl. This involves new techniques for scalability and for quality control in the presence of very noisy web contents (see Sections 4.1, 4.2 and 4.5).
- We extended the data model to include a new dimension of typicality scores, computed by a supervised regression model that considers assertion facets (see Section 4.6).
- We improved our techniques for canonicalizing assertions, added automatic mapping onto the schema of ConceptNet, and updated the evaluation. This includes a new module for statement canonicalization, enhanced clustering by leveraging SentenceTransformers [15], and evaluation with GPT-3 [16] (see Sections 4.3, 4.4 and 5.3).

2 RELATED WORK

Commonsense knowledge bases (CSKBs). CSK acquisition has a long tradition in AI (e.g., [3], [17], [18], [19]). A few projects have constructed large-scale collections that are publicly available. ConceptNet [4] is the most prominent project on CSK acquisition. Relying mostly on human crowdsourcing, it contains highly salient information for a small number of pre-specified predicates (isa/type, part-whole, used for, capable of, location of, plus lexical relations such as synonymy, etymology, derived terms etc.), and this CSKB is most widely used. However, it has limited coverage on many concepts, and its ranking of assertions, based on the number of crowdsourcing inputs, is very sparse and unable to discriminate salient properties against atypical or exotic ones (e.g., listing trees, gardens and the bible as locations of snakes, with similar scores). ConceptNet does not properly disambiguate concepts, leading to incorrect assertion chains like $\langle \textit{elephant}, \textit{hasPart}, \textit{trunk} \rangle$; $\langle \textit{trunk}, \textit{locationOf}, \textit{spare tire} \rangle$.

WebChild [5], TupleKB [6] and Quasimodo [7] devised fully automated methods for CSKB construction. They use judiciously selected text corpora (incl. book n-grams, image tags, QA forums) to extract large amounts of SPO triples. WebChild builds on hand-crafted extraction patterns, and TupleKB and Quasimodo rely on open information extraction (OpenIE) with subsequent cleaning. All three are limited to SPO triples.

ATOMIC [12], [20] and ASER [21] are recent projects on event-centered CSKB construction. While ATOMIC is entirely based on large-scale human compilation, assertions in ASER were extracted automatically from large text corpora, based on dependency patterns. Both resources use fixed relations (such as $x\textit{Need}$, $x\textit{Want}$, $\textit{isAfter}$ or $\textit{isBefore}$ in ATOMIC,

and *Reason, Result, Precedence* or *Succession* in ASER) and also are limited to triples. More recently, TransOMCS [22] mined concept-centric triples from ASER. TransOMCS used ConceptNet assertions as seed samples for pattern extraction, hence has the same set of pre-specified predicates as ConceptNet.

Another notable CSKB project is GenericsKB, which drops attempts at structuring assertions, and instead focuses on collecting generic sentences on a per-subject basis [23].

With ASCENT [14], [24] we developed an open information extraction methodology that captures semantic facets and multi-word compounds as subjects, enabling finer-grained knowledge representation, as well as avoiding common disambiguation errors. For knowledge source discovery, ASCENT used commercial search engine APIs. This choice ensured precision, but reduced recall, and strictly limited the scalability on monetary grounds. The present ASCENT++ approach overcomes this limitation by using a huge web crawl instead, with input several orders of magnitude larger.

Taxonomy and meronymy induction. The organization of concepts in terms of subclass and part-whole relationships, termed hypernymy and meronymy, has received great attention in NLP and web mining (e.g., [25], [26], [27], [28], [29], [30], [31], [32]). The hand-crafted WordNet lexicon [33] organizes over 100k synonym sets with respect to these relationships, although meronymy is sparsely populated.

Recent methods for large-scale taxonomy induction from web sources include WebIsADB [32], [34] building on Hearst patterns and other techniques, and the industrial GIANT ontology [35] based on neural learning from user-action logs and other sources.

Meronymy induction at large scale has been addressed by [23], [36], [37] with pre-specified and automatically learned patterns for refined relations like physical-part-of, member-of and substance-of.

Our approach includes relations of both kinds, by extracting knowledge about salient subgroups and aspects of subjects. In contrast to typical taxonomies and part-whole collections, our subgroups include many multi-word phrases: composite noun phrases (e.g., “circus lion”, “lion pride”) and adjectival and verbal phrases (e.g., “male lion”, “roaring lion”). Aspects cover additional refinements of subjects that do not fall under taxonomy or meronymy (e.g., “lion habitat” or “lion’s prey”).

Expressive knowledge representation and extraction. Modalities such as *always, often, rarely, never* have a long tradition in AI research (e.g., [38]), based on various kinds of modal logics or semantic frame representations, and semantic web formalisms can capture context using e.g., RDF* or reification [39]. While such expressive knowledge representations have been around for decades, there has hardly been any work that populated KBs with such refined models, notable exceptions being the Knext project [40] at small scale, and OntoSenticNet [41] with focus on affective valence annotations.

Other projects have pursued different kinds of contextualizations for CSK extraction, notably [42], which scored natural language sentences on an ordinal scale covering the spectrum *very likely, likely, plausible, technically possible* and *impossible*, [43] with probabilistic scores, and the Dice

project [44] which ranked assertions along the dimensions of plausibility, typicality and saliency.

Semantic role labelling (SRL) is a representation and methodology where sentences are mapped onto frames (often for certain types of events) and respective slots (e.g., agent, participant, instrument) are filled with values extracted from the input text [45], [46], [47]. Recently, this paradigm has been extended towards facet-based open information extraction, where extracted tuples are qualified with semantic facets like location and mode [48], [49]. ASCENT and ASCENT++ have built on this general approach, but extended it in various ways geared for the case of CSK: focusing on specifically relevant facets, refining subjects by subgroups and aspects, and aiming to reconcile precision and coverage for concepts as target subjects.

Pre-trained language models and CSK. Pre-trained language models (LMs) like BERT [50] and GPT variants [16], [51] have revolutionized NLP, with great advances on many tasks including those related to commonsense knowledge. For instance, the LAMA probe [52] aimed to extract commonsense knowledge from masked language models via a cloze-style QA task. COMET [53] is an autoregressive language model fine-tuned on existing CSK resources (e.g., ConceptNet and ATOMIC) that is used to predict objects for given subject-predicate pairs, and there are several related works [54], [55]. Nonetheless, as we show in Section 5.2, LM-based prediction tails extractive CSKB construction, once a comparable input corpus is used. In Section 5.3 we also show that explicit CSK can boost the performance of vanilla LMs in fill-in-the-blank tasks.

3 MODEL AND ARCHITECTURE

3.1 Knowledge model

Existing CSKBs typically follow a triple-based data model, where subjects are linked via predicate phrases to object words or phrases. Typical examples, from ConceptNet, are $\langle bus, UsedFor, travel \rangle$ and $\langle bus, UsedFor, not\ taking\ the\ subway \rangle$.

Few projects [5], [6] have attempted to sharpen such assertions by word sense disambiguation (WSD) [56], distinguishing, for example, buses on the road from computer buses. Likewise, only a few projects [7], [42], [44], [57], have tried to identify salient assertions against correct ones that are unspecific, atypical, or even misleading (e.g., buses used for avoiding the subway or used for enjoying the scenery). We extend this prevalent paradigm in two significant ways.

Expressive subjects. CSK acquisition starts by collecting assertions for target subjects, which are usually single nouns. This has two handicaps: 1) it conflates different meanings for the same word, and 2) it misses out on refinements and variants of word senses. While some projects [5], [6] tried to use word sense disambiguation (WSD) to overcome the first issue, they were inherently limited by the underlying word-sense lexicons (WordNet or Wiktionary). Indeed, they mainly restrict themselves to single nouns. For example, phrases like “city bus” or “tourist bus” are not present.

Our approach to rectify this problem is twofold:

- First, when discovering source documents for the target subject, we compare the documents with its reference Wikipedia article, and then only retain documents with

better similarity. This way, we can disentangle different senses, for example, of “bus” as in “public transport” and “network topology” themes.

- Second, when extracting candidates for assertions from the source documents, we capture also multi-word phrases as candidates for refined subjects, such as “school bus”, “city bus”, “tourist bus”, “circus elephant”, “elephant cow”, “domesticated elephant”, etc. This way, we can acquire *isa*-like refinements, to create *subgroups* of broader subjects, and also other kinds of *aspects* that are relevant to the general concept. An example for the latter would be “bus driver” or, for the target subject “elephant”, phrases such as “elephant tusk”, “elephant habitat” or “elephant keeper”.

Our notion of *subgroups* can be thought of as an inverse *isa* relation. It goes beyond traditional taxonomies by better coverage of multi-word composites (e.g., “circus elephant”). This allows us to improve the representation of specialized assertions such as $\langle \text{circus elephants, catch, balls} \rangle$.

Our notion of *aspects* includes part-whole relations (PartOf, MemberOf, SubstanceOf) [27], [36], [37], [58], but also further aspects that do not fall under the themes of hypernymy or meronymy. Examples are “elephant habitat” or “bus accident”. Note that, unlike single nouns, such compound phrases are rarely ambiguous, so we have crisp concepts without the need for explicit WSD.

Semantic facets. For CSK, assertion validity often depends on specific temporal and spatial circumstances, e.g., elephants scare away lions only in Africa or bathe in rivers only during daytime. Furthermore, assertions often become crisper by contextualization in terms of causes/effects and instruments (e.g., children ride the bus ... to go to school, circus elephants catch balls ... with their trunks).

To incorporate such information into an expressive model, we choose to contextualize subject-predicate-object triples with semantic *facets*. To this end, we build on ideas from research on semantic role labeling (SRL) [45], [46], [47]. This line of research has been initially devised to fill hand-crafted frames (e.g., purchase) with values for frame-specific roles (e.g., buyer, goods, price, etc.). We start with a set of 35 labels proposed in [48], a combination of those in the Illinois Curator SRL [46], and 22 hand-crafted ones derived from an analysis of semantic roles of prepositions in Wiktionary. As many of these are very special, we condense them into eight widely useful roles that are of relevance for CSK: four that qualify the validity of assertions (*degree, location, temporal, other-quality*), and four that capture other dimensions of context (*cause, manner, purpose, transitive objects*).

These design considerations lead us to the following knowledge model.

Definition [Commonsense Assertion]:

Let C_0 be a set of primary concepts of interest, which could be manually defined or taken from a dictionary.

Subjects for assertions include all $s_0 \in C_0$ as well as judiciously selected multi-word phrases that contain some s_0 .

Subjects are interrelated by *subgroup* and *aspect* relations: each s_0 can be refined by a set of subgroup subjects denoted $sg(s_0)$, and by a set of aspect subjects denoted

$asp(s_0)$. The overall set of subjects is $C := C_0 \cup sg_{C_0} \cup asp_{C_0}$.

A commonsense assertion for $s \in C$ is a quadruple $\langle s, p, o, F \rangle$ with single-noun or noun-phrase subject s , short phrases for predicate p and object o and a set F of semantic facets. Each facet $(k, v) \in F$ is a key-value pair with one of eight possible keys k and a short phrase as v . Note that a single assertion can have multiple key-value pairs with the same key (e.g., different spatial phrases). \square

3.2 Extraction architecture

Design considerations. CSK collection has three major design points: (i) the choice of sources, (ii) the choice of the extraction techniques, and (iii) the choice of cleaning or consolidating the extracting candidate assertions.

As *sources*, most prior works carefully selected high-quality input sources, including book n-grams [5], concept definitions in encyclopedic sources, and school text corpora about science [59]. These are often a limiting factor in the KB coverage. Moreover, even seemingly clean texts like book n-grams come with a surprisingly high level of noise and bias (cf. [60]). Focused queries for retrieving relevant web pages were used by [6], but the query formulations required non-negligible effort. Query auto-completion and question-answering forums were tapped by Quasimodo [7]. While this gave access to highly salient assertions, it was, at the same time, adversely affected by heavily biased and sensational contents (e.g., search-engine auto-completion for “snakes eat” suggesting “... themselves” and “... children”).

Search engines were also used in ASCENT [14] in combination with refined search queries that use hypernyms for disambiguation. Then, each retrieved web page was compared with a reference document of its corresponding subject to make sure that the web page was not off-topic. This helped ASCENT to collect very high-quality documents for various subjects such as animals (e.g., there are lots of web pages exclusively talking about animal facts). However, the method did not work well on subjects for which named instances are generally more prominent than the general concepts, such as “university” or “car”. For instance, while for “elephant”, the method extracted 508 assertions of frequency greater than one, for “university” it found a mere 12. ASCENT++ rectifies this by directly using a huge corpus as extraction source. This approach was also taken by ASER [21] and its derivative, TransOMCS [22], yet these approaches are recall-oriented and lack appropriate consolidation.

For the *extraction techniques*, choices range from co-occurrence- and pattern-based methods (e.g., [61]) and open information extraction (e.g., [6], [7]) to supervised learning for classification and sequence tagging. Co-occurrence works well for a few pre-specified, clearly distinguished predicates using distant seeds. Supervised extractors require training data for each predicate and thus have the same limitation. Recent approaches, therefore, prefer OpenIE techniques, and our extractors follow this trend, too.

For *knowledge consolidation*, early approaches kept all assertions from the ingest process (e.g., crowdsourcing [4]), whereas recent projects employed supervised classifiers or rankers for cleaning [6], [7], [42], [44], and also limited forms of clustering [6], [7] for canonicalization (taming semantic

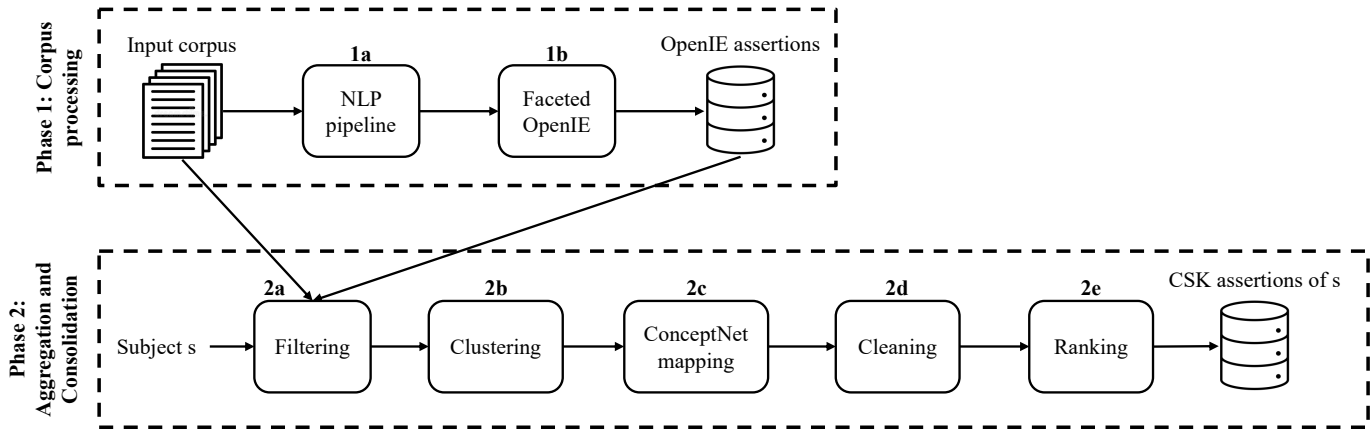


Fig. 1: Architecture of the ASCENT++ system.

redundancy). In ASCENT, we used contextual language models to cluster assertions of the same meanings and reinforce the frequency signals of those assertions. In the present ASCENT++, we apply the same technique with a state-of-the-art model for sentence embeddings which is the SentenceTransformers model [15]. Furthermore, after clustering, we do a mapping from our open-schema CSKB to the well-established ConceptNet schema as it is favorable to many researchers (e.g., [9], [20], [62]). Finally, a heuristic-based cleaning approach is applied to eliminate other remaining noise in the resulting KB.

Approach. The new ASCENT++ methodology operates in two phases (illustrated in Figure 1):

1. *Corpus processing*
 - 1a. *NLP pipeline*: Running NLP pipeline for the input corpus to get NLP features of all sentences, in particular part-of-speech tags and dependency trees.
 - 1b. *Faceted OpenIE*: Running the ASCENT OpenIE system to get faceted assertions from the processed sentences.
2. *Aggregation and consolidation*: in this phase, each subject is processed separately.
 - 2a. *Filtering*: Performing a series of document and assertion filtering to get relevant and high-quality assertions for a given subject.
 - 2b. *Clustering* of retained assertions based on sentence embeddings.
 - 2c. *Mapping* from open assertions into ConceptNet schema.
 - 2d. *Cleaning* based on heuristics and a dictionary of unwanted assertions.
 - 2e. *Ranking* of assertions: Annotating assertions with complementary scores for typicality and saliency.

Both phases treat each document or subject independently, thus can be highly parallelized (see also Section 5.1). In the following, we discuss Phase 1, and each of the subphases 2a-2e in a separate subsection.

4 METHODOLOGY

4.1 Corpus processing (Phase 1)

In ASCENT we retrieved 500 web pages returned by the Bing Search API using hand-crafted search queries for each subject. Usually, not all of those pages would be available for the extraction steps due to crawling errors and another document filtering step in the ASCENT pipeline. Thus, although this method gathered very high-quality sources, we had to sacrifice recall for several subjects.

To address these limitations of the web-search component in ASCENT, we drop it, and instead, ASCENT++ can use any English corpus as a knowledge source. This enables the system to tap texts on several orders of magnitude larger, making it easy to switch to any particular domain of interest.

Phase 1a: NLP pipeline. As we reuse the ASCENT OpenIE approach [14], the NLP pipeline consists of fundamental operations, including sentence splitting, tokenization, lemmatization, part-of-speech tagging, dependency parsing, and named entity recognition. The extractors will use all of these basic NLP features to output *faceted* OpenIE tuples.

Phase 1b: Faceted OpenIE. The ASCENT OpenIE approach was built upon StuffIE [48], a rule-based system capable of extracting triples and semantic facets from English sentences. The core ideas of the approach are to consider each verb as a candidate predicate of an assertion and then identify subjects, objects, and facets via grammatical relations, so-called dependency paths. Subjects must be connected to the candidate predicate through one of the subject-related dependency edges (n_{subj} , $n_{subjpass}$ and c_{subj}) or the adjectival clause edge (a_{cl}). Meanwhile, for objects, the respective edges include direct object (d_{obj}), indirect object (i_{obj}) and nominal modifier (n_{mod}). Next, the semantic facets are identified through the following complements to the selected verb: adverb modifiers, prepositional and clausal complements. We also extended the original set of rules in StuffIE to better deal with conjuncts and adverb facets which helps to identify significantly more assertions and facets and also improve the conciseness of the output tuples. For instance, given the sentence “Lynx are active during the evening and early morning”, our system can extract two assertions: $\langle lynx, are, active, TEMPORAL: during$

the evening) and $\langle \text{lynx, are, active, TEMPORAL: during early morning} \rangle$. Details on the dependency-pattern-based rules were previously described in Section 4.2 in [14].

Besides faceted OpenIE tuples, in ASCENT, we also proposed heuristics to extract fine-grained subjects, i.e., subgroups and aspects, of a given primary concept.

In terms of **subgroups**, they could be sub-species in the case of animals or refer to the primary subject in different states, such as “hunting cheetah” and “retired policeman”. For a primary subject s_0 , we collect all noun chunks ending with s_0 or any of its WordNet lemmas as potential candidates. Then, we cluster these terms based on word2vec [63] embeddings. In addition, we leverage WordNet to distinguish antonyms, with which the vector space embeddings typically struggle.

For **aspects**, given a primary subject s_0 , its aspects are extracted from noun chunks collected from two sources: (i) *possessive noun chunks*, for example, “elephant’s diet” and “their diet” (with resolution to “elephant”); (ii) $\langle s, p, \textit{noun chunk} \rangle$ triples where s is equal s_0 or any of its WordNet lemmas, and p is one of the following verb phrases: “have”, “contain”, “be assembled of” or “be composed of”.

The extracted subgroups and aspects are also normalized and cleaned in order to avoid spurious extractions or overly specific terms (cf. Section 4.2 in [14]).

4.2 Filtering (Phase 2a)

While the earlier search-engine-based document retrieval in ASCENT helped for topical relevance (a core focus of search engines), using a general web corpus as an extraction source implies the presence of substantially more irrelevant content.

Given a subject s , one might first collect all OpenIE assertions whose subject is equal s as candidate assertions and then apply ranking or filtering techniques afterward. Similar approaches have been used in Quasimodo [7] and TupleKB [6]. Nevertheless, such post-hoc filtering misses out on broader context from the original documents. In ASCENT++, we thus employ some filters first, at the document level, to decide which documents to use for candidate extraction at all. We only extract statements for a primary subject s_0 and its subgroups and aspects from a document d , if it passes the following filters:

- 1) Document d is only used, if it contains between 3 and 40 assertions for s_0 . The rationale for filters in either direction is that if s_0 occurs too rarely, d is more likely off-topic. If s_0 occurs too often, then d may be a noisy document such as a machine-created shopping catalog or simply a crawling error.
- 2) Similar to ASCENT, we compute the similarity between d and the Wikipedia article of the subject s_0 . Document d will be retained only if the similarity is above a certain threshold. This way, we can deal with ambiguous subject terms like “lynx”, which can be either an animal, a constellation, or a web browser.

After the previous filters, we come up with a list of documents for the subject s_0 . Then, we collect all OpenIE assertions for s_0 and its subgroups and aspects from those documents. After aggregating the extracted assertions, we only retain those with a frequency of at least 3. That helps to remove noise and redundancy from the results.

4.3 Clustering (Phase 2b)

Natural language is rich in paraphrases. For example, “lion eats zebra” can also be written as “lion preys on zebra” or “lions feed on zebra”. Therefore, identifying and clustering such assertions is necessary to avoid redundancies and get better frequency signals for individual assertions.

In ASCENT, we fine-tuned a RoBERTa model [64] to classify if two given triples have the same meaning. The training data consists of 21k triple pairs sampled from the extractions. The model outputs a number between 0 and 1, indicating the degree of semantic similarity between the two given triples. These similarity scores were used in the hierarchical agglomerative clustering (HAC) algorithm to group similar triples together.

In our new ASCENT++ system, we replace that classification model with SentenceTransformers [15], a state-of-the-art architecture for sentence embeddings. First, given a triple, we concatenate its subject, predicate, and object. Next, we feed the whole string to SentenceTransformers to get its contextualized embeddings. Then, for each pair of triples, we compute the cosine similarity of their corresponding embeddings. These cosine similarity scores will be used as input of the HAC algorithm. The model we use in our experiment is `paraphrase-mpnet-base-v2`¹, which was trained on multiple paraphrase datasets (e.g., `altlex`, `msmarco-triplets`, `quora_duplicates`, etc.) and has high performances in various evaluation schemes².

For facet clustering, we use average word2vec embeddings and the HAC algorithm, same as the approach previously used in ASCENT.

4.4 ConceptNet mapping (Phase 2c)

There are two main schools for knowledge representation in CSKBs: those relying on open predicates and those using a fixed set. Each has its strengths and challenges in terms of expressivity, redundancy, and usability. To bridge the two, we provide our ASCENT++ KB in two variants: with open assertions and with canonicalized predicates. Open information extraction supplies the former; the present module normalizes open assertions into fixed relations.

Our fixed schema of choice is the established ConceptNet schema [4], from which we use the following 19 commonsense relations: *AtLocation*, *CapableOf*, *Causes*, *CreatedBy*, *DefinedAs*, *Desires*, *HasA*, *HasPrerequisite*, *HasProperty*, *HasSubevent*, *IsA*, *MadeOf*, *MotivatedByGoal*, *PartOf*, *ReceivesAction*, *RelatedTo*, *SimilarTo*, *SymbolOf* and *UsedFor*.

Mapping open triples to a fixed schema raises several challenges. In the most straightforward case, subject and object from the open assertion can remain unchanged, and we only need to pick one of the fixed relations. For example, $\langle \textit{fish, lives in, water} \rangle$ can be mapped to $\langle \textit{fish, LocatedAt, water} \rangle$. In other cases, part of relation and object can be moved, e.g., $\langle \textit{handle, is, a part of a door} \rangle$ can be mapped to $\langle \textit{handle, PartOf, door} \rangle$. In other cases, part or all of the predicate is in the object, like in $\langle \textit{elephant, eats, grass} \rangle$ that can be mapped to $\langle \textit{elephant, CapableOf, eat grass} \rangle$.

1. <https://huggingface.co/sentence-transformers/paraphrase-mpnet-base-v2>

2. https://sbert.net/docs/pretrained_models.html

Our normalization method consists of two steps: (i) first, we use a multi-class classifier to predict a fixed predicate for each open triple; (ii) second, we use a list of crafted rules to modify the object so that the new triple preserves the meaning of the original triple.

For the classification model, we fine-tune a RoBERTa model to predict one of the ConceptNet predicates given an input in the following format: [CLS] S [SEP] P O, whereas S, P and O are the subject, predicate and object of the open triple. The contextualized embeddings of the [CLS] token will be fed to a fully connected layer to get the prediction. The training data for this model is constructed from ConceptNet triples. Specifically, for each ConceptNet predicate, we manually compiled a list of 1-6 open predicates with similar meanings. For instance, *UsedFor* is aligned with “be used for”, meanwhile *CapableOf* is aligned with “be capable of”, “be able to”, “can”, “could” and an empty string. This way, we can automatically generate 1.2M training examples.

After fine-tuning the LM on that generated data, we observed that it is sometimes confused between the three predicates, *IsA*, *HasProperty* and *ReceivesAction* which are all aligned to the open predicate “be”. For this particular case, we have a post-processing step to modify the predicted predicate: we only assign *HasProperty* to objects which are adjective phrases, *ReceivesAction* to objects which are verb phrases in passive form, and the rest are assigned to *IsA*. The *IsA* relation still contains considerable noise, hence, later in the cleaning phase (cf. Section 4.5), we introduce heuristics to get high-quality assertions of this type.

Once we get a predicted fixed-schema predicate, we have to produce an object for the normalized triple. The most common case when the object needs modifications is when the predicted predicate is *CapableOf*. In this case, if the open predicate is neither “be capable of”, “be able to”, “can” nor “could”, the new object will be the concatenation of the original predicate and object. In some cases, a part of the original object must be cut out as it overlaps with the ConceptNet predicate. Those predicates include *PartOf* and *SymbolOf*, for example, open triples corresponding to those predicates usually look like $\langle \textit{lion}, \textit{is}, \textit{part of a pride} \rangle$ or $\langle \textit{lion}, \textit{is}, \textit{symbol of strength} \rangle$ which should be translated to $\langle \textit{lion}, \textit{PartOf}, \textit{a pride} \rangle$ and $\langle \textit{lion}, \textit{SymbolOf}, \textit{strength} \rangle$ respectively. Other rules deal with other predicates such as *Desires*, *HasProperty*, *IsA*, *ReceivesAction* and *UsedFor*. In total, we have compiled 7 rules for object modification. If a triple does not fall into one of these rules, we preserve its original object.

4.5 Cleaning (Phase 2d)

Noise can come from various sources, including OpenIE errors, too specific or too general statements, nonsensical assertions, schema normalization errors, etc. Unlike supervised classifiers employed in other projects, our cleaning module is rule-based and thus highly scrutable. ASCENT++’s cleaning module consists of the following heuristics:

- First, we compute the perplexity of all triples. To do so, we first turn the triples into sentences and then use an autoregressive model (GPT-2) to compute the perplexity of the sentences. Only triples with medium-to-low perplexity will be retained (in our experiments, we retained triples with perplexity less than 500).

- The *IsA* triples produced by OpenIE are rich but rather noisy. Extracting *IsA* relations is a well-established research theme in entity typing and taxonomy construction and has already reached high-quality results. Therefore, we use the following cleaning algorithm. For each subject, we take the list of its *IsA* relations in ConceptNet and extract all head nouns of the objects. For example, from the ConceptNet triple $\langle \textit{elephant}, \textit{IsA}, \textit{placental mammal} \rangle$, we extract the head noun “mammal”. Then, in ASCENT++, we only retain the *IsA* assertions in which objects contain one of the trustworthy head nouns extracted earlier. This way, we not only get high-precision assertions but also better recall than ConceptNet. If a subject does not occur in ConceptNet, we remove all of its extracted *IsA* assertions, as in our observation, there are usually more noisy *IsA* assertions than valuable ones in the original extraction set. Note that, while ConceptNet is our source of choice here, any other existing resources that provide high-quality *IsA* relations can be used instead, for example, WordNet [33] or BabelNet [65].
- Then, we manually constructed a dictionary of unwanted objects from our observations. For example, we removed URLs, pronouns, numbers, only stop-words, too general/specific phrases such as “make sure” or “this case”, and vague predicate-object pairs (e.g., *MadeOf*-“part”, *HasProperty*-“available”). We also eliminate ethnicity- and religion-related assertions to avoid potentially critical biases [66].
- Finally, we only keep the 1,000 most frequent assertions per subject. For each assertion, we only keep its three most frequent facets. Cutting the tail this way improves the precision significantly but only minorly affects recall.

4.6 Ranking (Phase 2e)

Existing resources mainly provide unidimensional rankings of their assertions, by either *frequency* (e.g., ConceptNet, ASCENT), or supervised models trained to predict *plausibility/typicality* (TupleKB, Quasimodo, TransOMCS). These two dimensions are quite different, though, and we consider it important to differentiate their semantics.

Typicality states that an assertion holds for most instances of a concept. For example, elephants using their trunk is typical, whereas elephants drinking milk holds only for baby elephants. *Saliency* refers to the human perspective of whether an assertion is associated with a concept by most humans more or less on first thought. For example, “elephants have trunks” is salient, whereas “elephants pass by zebras” is not.

In ASCENT++, we make both dimensions first-class citizens of the CSKB and annotate each assertion with scores for both.

Saliency ranking. For saliency, we simply rely on the reporting frequency of the assertions, which approximates very well how prominent an assertion is. We transform raw frequencies to a normalized log-scaled frequency as follows:

$$\textit{saliency}(spo) = \frac{\log(\textit{count}(spo)) - \log(\textit{min_count}(s))}{\log(\textit{max_count}(s)) - \log(\textit{min_count}(s))} \quad (1)$$

Frequency adverb	Subject quantifier	Score
always	all, every	1.0
typically, mostly, mainly	most	0.9
usually, normally, regularly, frequently, commonly		0.8
	many	0.7
often		0.6
	some	0.5
sometimes		0.4
occasionally	few	0.3
		0.2
hardy, rarely		0.1
	no, none	0.0

TABLE 1: Frequency modifiers and assigned scores.

where $count(spo)$ is the raw frequency of the triple spo , $min_count(s)$ is the minimal frequency of a triple whose subject is s , and $max_count(s)$ is the maximal one.

Typicality ranking. Typicality is the most challenging dimension. Although reporting frequency correlates with typicality moderately, reporting bias in texts [60] means that sensational statements may be grossly overreported and commonalities underrepresented. Our qualitative facets (cf. Section 3.1) give us a unique handle to obtain further insights on typicality.

We use a linear regression model on three features:

- *Modifier score.* This feature is based on adverbs and quantifiers in facets and subjects. Specifically, we assign every frequency-related modifier a specific numeric score (see Table 1), then average all scores in each assertion cluster. We consider two types of modifiers: adverbs (e.g., “always”, “often”) that appear in semantic facets, and subject quantifiers such as “all”, “few” or “some”. In the absence of any modifier, we assign a default score of 0.5.
- *Neutrality.* First, we use a sentiment analysis model³ to compute the probability of a given sentence being positive, neutral, or negative. Then, for each assertion, we consider the average polarity over all of its source sentences as its polarity. The value of this feature is 1 if the assertion is classified as neutral. Otherwise, a value of zero reflects a polarized assertion.
- *Normalized frequency.* The value for this feature is computed as in Equation 1.

5 EXPERIMENTS

The evaluation of ASCENT++ is centered on three research questions:

- **RQ1:** Is the resulting CSKB of higher quality than existing resources?
- **RQ2:** Does (structured) CSK help in extrinsic use cases?
- **RQ3:** What is the quality and extrinsic value of facets?

We first present the implementation of ASCENT++, then discuss each of these research questions in its own subsection.

3. <https://huggingface.co/cardiffnlp/twitter-roberta-base-sentiment>

5.1 Implementation

We executed the new ASCENT++ pipeline for all primary subjects in ASCENT (i.e., 10K popular subjects taken from ConceptNet), and for each primary subject, we also took its top 10 most frequent aspects and top 10 most frequent subgroups previously extracted by ASCENT as high-quality fine-grained subjects. In Table 2, we provide the size of our input corpus and the number of retained assertions at different stages in the pipeline.

Our input corpus of choice is the C4 corpus, created to train the T5 text-to-text Transformers model [13]. C4 was created by cleaning the Common Crawl’s web crawl corpus. The version we use consists of 365M English articles. Each comes with its text, URL, and crawling timestamp. The extraction phase in ASCENT++ is executed once, independent of any choice of subjects. We use SpaCy⁴, a popular Python-based NLP library, as our NLP pipeline. First, we ran SpaCy on the 365M documents from C4 in parallel on a cluster of 6,400 CPU cores (AMD EPYC 7702 64-core processors), which took 1.5 days to complete. Then, it took 20h for the OpenIE system to digest all 356M processed documents on the same CPU cluster. The result of this step is a set of 8B OpenIE assertions.

Next, the filtering process took around 2 hours and resulted in 165M assertions for the selected subjects, which account for 15M unique triples.

The clustering process, including precomputing embeddings and running the HAC algorithm, took around 10 hours. Embeddings were computed on a cluster of up to 150 GPUs (NVIDIA Quadro RTX 8000 GPUs) which took less than 2 hours. The HAC algorithm was run on the same CPU cluster, giving 7.5M clustered assertions. On average, two assertions form one cluster.

Next, the ConceptNet mapping took around 30 minutes on the same GPU cluster. The rule-based cleaning step is the lightest component, taking less than 10 minutes on a personal laptop. It resulted in 2M assertions in the final ASCENT++ CSKB, for which important statistics are provided in Table 3. In particular, there are about 10K subgroups and 5.8K aspects. We have 1.6M assertions for primary subjects, 80K assertions for subgroups, and 323K assertions for aspects. The total number of semantic facets is 2.3M. Hence, each assertion has more than one facet on average.

For typicality scoring, we trained a regression model on 500 manually-annotated examples. The resulting formula is: $typicality = 0.324 \times modifier_score + 0.428 \times frequency + 0.088 \times neutrality$. The ranking module took less than 2 hours to complete on the same CPU and GPU clusters.

We show examples highlighting markedly different ranks by typicality and saliency in Table 4.

5.2 Intrinsic evaluation

To investigate RQ1, we instantiate *quality* with the standard notions of precision and recall, splitting precision further up into the dimensions of *typicality* and *saliency*, measuring this way the degree of truth, and the degree of relevance of assertions (cf. [7]).

4. <https://spacy.io/>

Input documents	All OpenIE assertions	Filtered assertions	Unique triples	Clustered assertions	Clean assertions
365M	8B	165M	15M	7.5M	2M

TABLE 2: Number of input documents and assertions at different stages in the pipeline.

Subject type	#s	#spo	#facets
Primary	8,067	1,651,455	1,975,385
Subgroup	10,191	80,176	62,581
Aspect	5,843	323,257	312,004
All	24,101	2,054,888	2,349,970

TABLE 3: Statistics on different subject types in ASCENT++.

ASCENT++ triple	Typ. rank	Sal. rank
<i><elephant, CapableOf, eat grass></i>	2	8
<i><elephant, CapableOf, eat fruit></i>	3	12
<i><elephant, CapableOf, destroy crop></i>	26	1
<i><car, AtLocation, accident></i>	9	1
<i><lawyer, CapableOf, understand the law></i>	7	50

TABLE 4: Examples on how statement ranks based on typicality (typ.) and saliency (sal.) differ.

Compared resources. We compare ASCENT++ with its predecessor ASCENT [14], as well as five other prominent resources:

- 1) ConceptNet [4],
- 2) TransOMCS [22],
- 3) TupleKB [6],
- 4) Quasimodo [7],
- 5) COMET [20] (with reservations, see below).

The ATOMIC [12] and ASER [21] projects do not qualify for comparison, as they do not contain concept-centric assertions.

Assertion precision. Unlike for encyclopedic knowledge (“The Lion King” was either produced by Disney, or it was not), the precision of CSK is generally not a binary concept, calling for more refined evaluation metrics. Following the Quasimodo project [7], we assessed the *typicality* and *saliency* of triples. Given a CSK triple, we asked annotators on Amazon MTurk to rate the triple along a 4-point Likert scale on each of the two dimensions. We present the crowdsourcing questions and answer options in Table 5. Note that we followed the exact same scheme introduced in COMET [20] for typicality evaluation. Hence, besides the four options in Table 5 there is also an option for **No judgment** if the concepts presented in the assertion are too unfamiliar for the assigned worker. We performed two separated MTurk tasks on the two dimensions. For each task, we randomly sampled 500 triples among 200 prominent subjects in four common domains: animal, occupation, engineering, and geography. For the saliency task, the sampling pool was limited to the top 10 assertions per subject.

Each MTurk task contained five CSK triples and was assigned to three different workers. Following [20], we also used human-readable language forms for triples in the fixed-schema CSKBs (i.e., ConceptNet, TransOMCS and COMET). For ASCENT++, we used the open triples as the prompt display. Any triples that receive the first two labels (cf.

Table 5) are ranked as *Typical/Salient*, and the last two labels as *Untypical/Unsalient*. The final judgment is based on a majority vote. Annotation quality was ensured by requiring workers to be Master workers with an all-time approval rate of over 90%. The inter-rater agreement on the three labels measured by Fleiss’ κ is 0.33 (i.e., fair agreement [67]).

We report the precision evaluation results in the left part of Table 6. The results are consistent with those presented in the ASCENT paper [14]. Among automatically-constructed CSKBs, ASCENT has the most salient triples by a large margin. In terms of typicality, ASCENT also has a competitive performance. Now, the new ASCENT++ KB goes even further. By switching to a massive amount of texts as the new knowledge source and using diverse filtering and cleaning techniques, we saw an increment of 8.8 and 9.2 percentage points over ASCENT on triple saliency and typicality, respectively.

We also compared our KB with a generative CSKB construction method, COMET [53]. COMET does not come with a resource, hence we had to generate assertions ourselves, and as there is no obvious stop criterion, we only evaluated precision of top assertions, but could not evaluate its recall. The model we used is the provided BART model⁵ which was trained on ATOMIC 2020 (which also includes ConceptNet assertions). For each pair of subject and predicate, we asked COMET to predict five objects. We used the same sampling processes and MTurk templates described above for typicality and saliency evaluation. The evaluation results of COMET are also included in Table 6. ASCENT++ clearly has a better performance than its generative counterpart, even though both have seen comparable amounts of texts. Some examples of the top assertions in ASCENT++ and the ones generated by COMET are shown in Table 7. Although COMET has the flexibility of generating objects to any given subject-predicate pair, it still makes many incorrect predictions and produces notable redundancies. On the other hand, ASCENT++, which collected OpenIE assertions from several sources and aggregated them through various steps such as filtering, clustering, and cleaning, produces more correct and complementary assertions.

Assertion recall. Evaluating recall requires a notion of ground truth. We used the CSLB property norm dataset [68] which consists of short human-written sentences about salient properties of general concepts. There are 22.6K sentences about 638 concepts in the dataset. This replaced the much smaller crowd-sourced dataset of 2.4K sentences used in the Quasimodo project [7].

As there are always different expressions of a CSK assertion in natural language, we allow soft matching between triples in CSKBs and sentences in the ground-truth CSLB dataset. To do so, we first transform predicates in the CSKBs into the schema of CSLB (i.e., only 3 predicates: “is”, “has”, “does”). Then for each CSLB sentence, we find the most

5. https://github.com/allenai/comet-atomic-2020/tree/master/models/comet_atomic2020_bart

Dim.	Question	Options
Typicality	Is this a correct assertion about <subject>?	Always/Often - the knowledge assertion presented is always or often true, Sometimes/Likely - it is sometimes or likely true, Farfetched/Never - it is false or farfetched at best, Invalid - it is invalid or makes no sense
Saliency	Imagine you have 2 minutes time to explain a kid about <subject>, would you mention the following information?	Absolutely - the information is very interesting/important for that concept, Probably - it is quite good to know, Maybe not - it is not interesting or too obvious/uninteresting/boring, Definitely not - it is completely irrelevant/unimportant/wrong or makes no sense.

TABLE 5: Crowdsourcing question templates for typicality and saliency evaluation of CSK assertions.

CSK resource	Precision (%)				Recall (%)						Size
	Saliency@10		Typicality (all)		Top-100 assertions/subject			All assertions			
	Salient	Unsalient	Typical	Untypical	t = 0.96	t = 0.98	t = 1.00	t = 0.96	t = 0.98	t = 1.00	
<i>Crowdsourced</i>											
ConceptNet [4]	79.2	20.4	96.0	3.6	5.29	3.39	1.10	5.47	3.53	1.13	0.5M
<i>Generative</i>											
COMET [20]	55.2	43.6	78.9	20.1	-	-	-	-	-	-	-
<i>Extractive</i>											
TransOMCS [22]	40.4	59.0	51.4	46.2	4.04	3.24	1.85	19.70	16.47	9.06	18.5M
TupleKB [6]	36.0	63.0	92.0	6.4	3.57	1.99	0.41	4.32	2.49	0.58	0.3M
Quasimodo [7]	38.8	60.2	67.8	31.0	11.05	9.47	5.17	21.87	19.36	10.88	6.3M
ASCENT [14]	60.0	38.6	79.2	15.8	9.60	7.02	3.17	17.09	12.62	5.82	8.6M
ASCENT++	68.8	30.8	88.4	8.8	12.70	10.70	5.90	17.54	14.64	7.95	2.0M

TABLE 6: Intrinsic evaluation results of different CSK resources.

Subject-Predicate	ASCENT++	COMET
elephant - CapableOf	- perform trick 🍑 - eat grass 🍑 - eat fruit 🍑 - become agitated - give ride 🍑	- climb tree - walk on land 🍑 - climb tree trunk - walk on tree - eat elephant
beer - Made(Up)Of	- hop 🍑 - water 🍑 - barley 🍑 - yeast 🍑 - grain 🍑	- beer - alcohol 🍑 - drunk - drinking - drink
laptop - UsedFor/ ObjectUse	- work 🍑 - gaming 🍑 - office work 🍑 - email 🍑 - social media 🍑	- browse the internet 🍑 - use as a coaster - play games on 🍑 - use as a weapon - browse the web 🍑

TABLE 7: Top-5 assertions of selected subject-predicate pairs in ASCENT++ and COMET.

similar triple in the target CSKB based on cosine similarity between their SentenceTransformers [15] embeddings. The found triple is considered a match with the given CSLB fact if the cosine similarity is greater than or equal to a predefined threshold t . The recall evaluation results are shown in the right part of Table 6, where we once show recall using the top-100 assertions per subject, and once all. We report results at three similarity thresholds: $t = 0.96$, $t = 0.98$ and $t = 1.0$.

When it comes to the top assertions, ASCENT++ outperforms all other CSKBs at all three thresholds. This affirms that the ranking approach used in ASCENT++ helps to pull out

essential assertions to the top better than other KBs. When there are no restrictions, TransOMCS and Quasimodo have higher recall than the others due to their massive size (18.5M for TransOMCS and 6.3M for Quasimodo, while it is 2M for ASCENT++). However, these sizes come with considerable precision problems.

Precision of ConceptNet mapping. To evaluate the ConceptNet mapping module, we manually annotate 100 random samples from the final ASCENT++ KB. For each sample, we mark it as a correct mapping if the fixed-schema triple preserves the meaning of the original triple. The obtained precision is 96%.

5.3 Extrinsic evaluation

To answer RQ2, we conduct a comprehensive evaluation of the contribution of commonsense knowledge to question answering (QA) via three different setups, all based on the idea of priming pre-trained LMs with context [16], [69], [70]:

- 1) In *masked prediction* (MP) [52], we ask language models to predict single tokens in generic sentences.
- 2) In *autoregressive LM-based QA* (AR), we provide questions, and let LMs generate arbitrary answer sentences.
- 3) In *span prediction* (SP), LMs select best answers from provided CSKB content [71].

We illustrate all settings in Table 8. In all settings, LMs are provided with a context in the form of assertions taken from competitor CSKBs. These setups are motivated by the observation that priming language models with context can significantly influence their predictions [69], [70]. Previous

works on language model priming mainly focused on evaluating retrieval strategies. In contrast, our comprehensive test suite focuses on the impact of utilizing different CSK resources while leaving the retrieval component constant.

Masked prediction is perhaps the best-researched problem, coming with the advantage of allowing automated evaluation, although automated evaluation may unfairly discount sensible alternative answers. Also, masked prediction is limited to single tokens. Autoregressive LM-based generations circumvent this restriction, although they necessitate human annotations and can be prone to evasive answers. They are thus well complemented by extractive answering schemes, limiting the language models’ abstraction abilities but providing the cleanest way to evaluate the context alone.

Models. Following standard usage, we use RoBERTa-large [64] for masked prediction, GPT-3 [16] for the generative setup, and ALBERT-xxlarge [71], fine-tuned on SQuAD 2.0, for span prediction.

Context retrieval method. Given a query, we use sentence embeddings to pull out relevant assertions from a CSKB. Using the SentenceTransformers model `msmarco-distilbert-base-v3`⁶, we compute embeddings of the given query and all lexicalized triples in the CSKBs. Then we use cosine similarity to select the top 5 most similar triples to the query as context.

Task construction. Previous work has generated masked sentences based on templates from ConceptNet triples [52]. However, the resulting sentences are often unnatural, following the idiosyncrasies of the ConceptNet data model. We therefore created a new dataset of natural commonsense sentences for *masked prediction* which is based on an independent human-annotated dataset, the aforementioned CSLB dataset [68], and consists of 19.6K masked sentences [14].

For the *generative* and *extractive* settings, we used the Google Search auto-completion functionality to collect commonsense questions about common subjects by feeding the API with 6 prefixes: “what/when/where are/do <subject>...”. That process returned 8,098 auto-completed queries. Next, we drew samples from the query set, then manually removed jokes and other noise (e.g., “where do cows go for entertainment”), obtaining 100 questions for evaluation.

Evaluation scheme. For commonsense topics, questions often have multiple valid answers. Additionally, given that answers in our generative and extractive QA settings are very open, creating an automated evaluation is difficult. We, therefore, use human judgments for evaluating all settings except masked prediction. Specifically, given a question and set of answers, we ask humans to assess each answer based on two dimensions, *correctness* and *informativeness*, each on a 4-point Likert scale from 0 (lowest) to 3 (highest). Three annotators evaluate each question in Amazon MTurk with the same qualification requirements as in Section 5.2. For evaluating masked prediction, we use the mean precision at k ($P@k$) metric, following [52].

Results. The evaluation results are shown in Table 9.

For *masked prediction*, all CSKBs contribute useful contexts that substantially improve the quality of LM responses. ASCENT++ along with the only manually-constructed KB, ConceptNet, statistically significantly outperforms all other KBs at every threshold k (all p-values of paired t-test below 0.05).

For *autoregressive LM-based QA*, note that here we use GPT-3 [16] instead of GPT-2 [51] which was used in the ASCENT paper [14]. GPT-3 has ten times more parameters than GPT-2 and performs much better than its predecessor in several natural language-related tasks. GPT-3 rarely gives evasive answers like GPT-2 (hence, we dropped the *guided generation* setting previously introduced in the ASCENT paper). Markedly, we find that GPT-3 performs on average better without any context, although in combination with ASCENT++, it still performs better than in combination with any other CSKB. However, more research is needed to design methods to pull relevant context from CSKBs and decide when to use it in LM-based QA and when to rely on the LM’s knowledge alone.

For *span prediction* where answers come directly from retrieved contexts, ASCENT++ also outperforms all other competitors. ASCENT++ obtained statistically significant gains over Quasimodo, TupleKB, and TransOMCS on both metrics, and over ConceptNet on correctness. This indicates that our ASCENT++ assertions have high quality compared to others.

5.4 Evaluation of facets

To answer RQ3, we evaluate facets both intrinsically and extrinsically.

For *intrinsic evaluation*, as there are no existing CSKBs with facets, we provide comparisons with two strong LM baselines, GPT-2 [51] and GPT-3 [16]. First, we randomly drew 300 assertions along with their top-1 facets from our KB. Next, we translate each statement into a sentence prefix and ask the two GPT models to fill in the remaining words to complete the sentence. For example, given the quadruple $\langle \textit{elephant, use, their trunks, PURPOSE: to suck up water} \rangle$, the sentence prefix will be “Elephants use their trunk to” and for this, GPT-2’s continuation is “to move around” while GPT-3’s is “to breath”. Then, each sentence prefix along with three answers (from ASCENT++, GPT-2 and GPT-3) were shown to a human annotator (without knowing the source of the answers) who determined whether the answers were correct/incorrect and informative/uninformative. The results are reported in Table 10. ASCENT++ achieves the correctness of 70.10% and informativeness of 54.15%, which are both significantly better than GPT-2. However, ASCENT++ and GPT-2 are both outperformed by GPT-3 which was trained on a much larger dataset (approximately 5 times larger than the C4 dataset).

For *extrinsic evaluation*, we reused the three question answering tasks from Section 5.3. The results are shown in Table 11. Expanding the ASCENT++ triples with facets gives a consistent improvement in four of five evaluation metrics (precision at one in MP, informativeness in AR and SP, and correctness in SP), with the biggest effect being observed for span prediction (8.6% and 9.3% relative improvements over the no-facet context in informativeness and correctness, respectively).

6. <https://www.sbert.net/docs/pretrained-models/msmarco-v3.html>

Setting	Input	Sample output
MP	Elephants eat [MASK]. [SEP] Elephants eat roots, grasses, fruit, and bark, and they eat a lot of these things.	everything (15.52%), trees (15.32%), plants (11.26%)
AR	Context: Elephants eat roots, grasses, fruit, and bark, and they eat a lot of these things. Question: What do elephants eat? Answer:	Elephants eat a variety of different foods.
SP	question="What do elephants eat?" context="Elephants eat roots, grasses, fruit, and bark, and they eat a lot of these things."	start=14, end=46, answer="roots, grasses, fruit, and bark"

TABLE 8: Examples of three QA settings (MP - masked prediction, AR - autoregressive LM-based QA, SP - span prediction). Sample output was given by RoBERTa (for MP), GPT-3 (for AR) and ALBERT (for SP).

Context	MP			AR		SP	
	P@1	P@5	P@10	C	I	C	I
No context	8.10	17.16	21.37	2.47	2.01	-	-
ConceptNet	14.41	27.08	32.16	2.22	1.70	1.74	1.52
TransOMCS	7.08	15.42	19.99	1.32	0.86	0.99	0.85
TupleKB	11.61	24.76	30.36	2.22	1.51	1.70	1.38
Quasimodo	12.11	22.75	27.71	2.03	1.51	1.75	1.44
ASCENT	11.95	24.70	29.70	2.25	1.76	1.88	1.60
ASCENT++	13.30	27.03	32.90	2.32	1.71	1.94	1.63

TABLE 9: QA evaluation results. Metrics: P@k - mean precision at k (%), C - correctness ([0, 3]), I - informativeness ([0, 3]).

	Correct	Informative
GPT-2 [51]	61.79	48.84
GPT-3 [16]	80.07	70.76
ASCENT++	70.10	54.15

TABLE 10: Assessment of ASCENT++ and LM-generated facets (%).

6 CONCLUSION

This paper presented ASCENT++, a methodology to extract and semantically organize advanced commonsense knowledge from large-scale web contents. Our refined knowledge representation allowed us to identify considerably more informative assertions, overcoming limitations of prior works. The techniques for filtering, aggregating and consolidating extracted tuples show that CSK extraction from broad web content is feasible at scale, with both high precision and high recall. Intrinsic and extrinsic evaluations confirmed that the resulting CSKB is a significant advance over existing CSK collections and provides an edge over recent language-model-based approaches. Code, data, and a web interface are accessible at <https://www.mpi-inf.mpg.de/ascentpp>.

REFERENCES

[1] J. McCarthy, *Programs with common sense*. RLE and MIT computation center, 1960.
[2] E. A. Feigenbaum, "Knowledge engineering," *Annals of the New York Academy of Sciences*, 1984.
[3] D. B. Lenat, "Cyc: A large-scale investment in knowledge infrastructure," *CACM*, 1995.

Context	MP		AR		SP	
	P@1	C	I	C	I	
Without facets	13.30	2.32	1.71	1.94	1.63	
With facets	13.38	2.26	1.79	2.12	1.77	

TABLE 11: Extrinsic evaluation of facets by mean precision at one (%), correctness C ([0-3]) and informativeness I ([0-3]).

- [4] R. Speer and C. Havasi, "ConceptNet 5: A large semantic network for relational knowledge," *Theory and Applications of Natural Language Processing*, 2012.
[5] N. Tandon, G. de Melo, F. M. Suchanek, and G. Weikum, "WebChild: harvesting and organizing commonsense knowledge from the web," in *WSDM*, 2014.
[6] B. D. Mishra, N. Tandon, and P. Clark, "Domain-targeted, high precision knowledge extraction," *TACL*, 2017.
[7] J. Romero, S. Razniewski, K. Pal, J. Z. Pan, A. Sakhadeo, and G. Weikum, "Commonsense properties from query logs and question answering forums," in *CIKM*, 2019.
[8] H. Lin, L. Sun, and X. Han, "Reasoning with heterogeneous knowledge for commonsense machine comprehension," in *EMNLP*, 2017.
[9] B. Y. Lin, X. Chen, J. Chen, and X. Ren, "Kagnet: Knowledge-aware graph networks for commonsense reasoning," in *EMNLP-IJCNLP*, 2019.
[10] J. Xia, C. Wu, and M. Yan, "Incorporating relation knowledge into commonsense reading comprehension with multi-task learning," in *CIKM*, 2019.
[11] F. Ilievski, P. Szekely, and D. Schwabe, "Commonsense knowledge in Wikidata," *Wikidata workshop*, 2020.
[12] M. Sap, R. LeBras, E. Allaway, C. Bhagavatula, N. Lourie, H. Rashkin, B. Roof, N. A. Smith, and Y. Choi, "Atomic: An atlas of machine commonsense for if-then reasoning," in *AAAI*, 2018.
[13] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *Journal of Machine Learning Research*, 2020.
[14] T.-P. Nguyen, S. Razniewski, and G. Weikum, "Advanced semantics for commonsense knowledge extraction," in *WWW*, 2021.
[15] N. Reimers and I. Gurevych, "Sentence-bert: Sentence embeddings using siamese bert-networks," in *EMNLP*, 2019.
[16] T. B. Brown *et al.*, "Language models are few-shot learners," in *NeurIPS*, 2020.
[17] P. Singh, T. Lin, E. T. Mueller, G. Lim, T. Perkins, and W. L. Zhu, "Open mind common sense: Knowledge acquisition from the general public," in *OTM Confederated International Conferences*, 2002.
[18] H. Liu and P. Singh, "Conceptnet—a practical commonsense reasoning tool-kit," *BT technology journal*, 2004.
[19] J. Gordon, B. V. Durme, and L. K. Schubert, "Learning from the web: Extracting general world knowledge from noisy text," in *AAAI Workshops*, 2010.
[20] J. D. Hwang, C. Bhagavatula, R. L. Bras, J. Da, K. Sakaguchi, A. Bosselut, and Y. Choi, "(Comet-)Atomic 2020: On symbolic and neural commonsense knowledge graphs," in *AAAI*, 2021.

- [21] H. Zhang, X. Liu, H. Pan, Y. Song, and C. W.-K. Leung, "Aser: A large-scale eventuality knowledge graph," in *WWW*, 2020.
- [22] H. Zhang, D. Khashabi, Y. Song, and D. Roth, "Transomcs: From linguistic graphs to commonsense knowledge," in *IJCAI*, 2020.
- [23] S. Bhakthavatsalam, C. Anastasiades, and P. Clark, "Genericskb: A knowledge base of generic statements," *arXiv:2005.00660*, 2020.
- [24] T.-P. Nguyen, S. Razniewski, and G. Weikum, "Inside ASCENT: Exploring a deep commonsense knowledge base and its usage in question answering," in *ACL: System Demonstrations*, 2021.
- [25] O. Etzioni, M. J. Cafarella, D. Downey, S. Kok, A. Popescu, T. Shaked, S. Soderland, D. S. Weld, and A. Yates, "Web-scale information extraction in knowitall: (preliminary results)," in *WWW*, 2004.
- [26] R. Snow, D. Jurafsky, and A. Y. Ng, "Semantic taxonomy induction from heterogeneous evidence," in *ACL*, 2006.
- [27] R. Girju, A. Badulescu, and D. I. Moldovan, "Automatic discovery of part-whole relations," *Comput. Linguistics*, 2006.
- [28] P. Pantel and M. Pennacchiotti, "Espresso: Leveraging generic patterns for automatically harvesting semantic relations," in *ACL*, 2006.
- [29] M. Pasca and B. V. Durme, "Weakly-supervised acquisition of open-domain classes and class attributes from web documents and query logs," in *ACL*, 2008.
- [30] S. P. Ponzetto and M. Strube, "Taxonomy induction based on a collaboratively built knowledge repository," *Artif. Intell.*, 2011.
- [31] W. Wu, H. Li, H. Wang, and K. Q. Zhu, "Probase: A probabilistic taxonomy for text understanding," in *SIGMOD*, 2012.
- [32] S. Hertling and H. Paulheim, "WebIsALOD: Providing hypernymy relations extracted from the web as linked open data," in *ISWC*, 2017.
- [33] G. A. Miller, "Wordnet: A lexical database for English," *CACM*, 1995.
- [34] J. Seitner, C. Bizer, K. Eckert, S. Faralli, R. Meusel, H. Paulheim, and S. P. Ponzetto, "A large database of hypernymy relations extracted from the web," in *REC*, 2016.
- [35] B. Liu, W. Guo, D. Niu, J. Luo, C. Wang, Z. Wen, and Y. Xu, "GIANT: scalable creation of a web-scale ontology," in *SIGMOD*, 2020.
- [36] N. Tandon, C. Hariman, J. Urbani, A. Rohrbach, M. Rohrbach, and G. Weikum, "Commonsense in parts: Mining part-whole relations from the web and image tags," in *AAAI*, 2016.
- [37] S. Bhakthavatsalam, K. Richardson, N. Tandon, and P. Clark, "Do dogs have whiskers? a new knowledge base of haspart relations," *arXiv:2006.07510*, 2020.
- [38] D. M. Gabbay, *Many-Dimensional Modal Logics: Theory and Applications*. Elsevier North Holland, 2003.
- [39] A. Hogan *et al.*, "Knowledge graphs," *ACM Computing Surveys (CSUR)*, 2021.
- [40] L. Schubert, "Can we derive general world knowledge from texts," in *HLT*, 2002.
- [41] M. Dragoni, S. Poria, and E. Cambria, "Ontosentinet: A commonsense ontology for sentiment analysis," *IEEE Intell. Syst.*, 2018.
- [42] S. Zhang, R. Rudinger, K. Duh, and B. Van Durme, "Ordinal common-sense inference," *TACL*, 2017.
- [43] T. Chen, Z. Jiang, A. Poliak, K. Sakaguchi, and B. Van Durme, "Uncertain natural language inference," in *ACL*, 2020.
- [44] Y. Chaliier, S. Razniewski, and G. Weikum, "Joint reasoning for multi-faceted commonsense knowledge," in *AKBC*, 2020.
- [45] M. Palmer, D. Gildea, and N. Xue, *Semantic Role Labeling*, ser. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers, 2010.
- [46] J. Clarke, V. Srikumar, M. Sammons, and D. Roth, "An NLP curator (or: How i learned to stop worrying and love NLP pipelines)," in *REC*, 2012.
- [47] G. Stanovsky, J. Michael, L. Zettlemoyer, and I. Dagan, "Supervised open information extraction," in *NAACL*, 2018.
- [48] R. E. Prasojo, M. Kacimi, and W. Nutt, "Stuffie: Semantic tagging of unlabeled facets using fine-grained information extraction," in *CIKM*, 2018.
- [49] M. Cetto, C. Niklaus, A. Freitas, and S. Handschuh, "Graphene: Semantically-linked propositions in open information extraction," in *COLING*, 2018.
- [50] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *NAACL*, 2019.
- [51] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [52] F. Petroni, T. Rocktäschel, P. Lewis, A. Bakhtin, Y. Wu, A. H. Miller, and S. Riedel, "Language models as knowledge bases?" in *EMNLP*, 2019.
- [53] A. Bosselut, H. Rashkin, M. Sap, C. Malaviya, A. Çelikyilmaz, and Y. Choi, "COMET: commonsense transformers for automatic knowledge graph construction," in *ACL*, 2019.
- [54] J. Davison, J. Feldman, and A. M. Rush, "Commonsense knowledge mining from pretrained models," in *EMNLP*, 2019.
- [55] Z. Bouraoui, J. Camacho-Collados, and S. Schockaert, "Inducing relational knowledge from bert," in *AAAI*, 2020.
- [56] R. Navigli, "Word sense disambiguation: A survey," *ACM Comput. Surv.*, 2009.
- [57] J. Gordon and L. K. Schubert, "Quantificational sharpening of commonsense knowledge," in *Commonsense Knowledge, AAAI Fall Symposium*, 2010.
- [58] V. Shwartz and C. Waterson, "Olive oil is made of olives, baby oil is made for babies: Interpreting noun compounds using paraphrases in a neural model," in *NAACL*, 2018.
- [59] P. Clark, I. Cowhey, O. Etzioni, T. Khot, A. Sabharwal, C. Schoenick, and O. Tafjord, "Think you have solved question answering? try ARC, the AI2 reasoning challenge," *arXiv:1803.05457*, 2018.
- [60] J. Gordon and B. V. Durme, "Reporting bias and knowledge acquisition," in *AKBC*, 2013.
- [61] Y. Elazar, A. Mahabal, D. Ramachandran, T. Bedrax-Weiss, and D. Roth, "How large are lions? inducing distributions over quantitative attributes," in *ACL*, 2019.
- [62] Y. Feng, X. Chen, B. Y. Lin, P. Wang, J. Yan, and X. Ren, "Scalable multi-hop relational reasoning for knowledge-aware question answering," in *EMNLP*, 2020.
- [63] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *NeurIPS*, 2013.
- [64] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv:1907.11692*, 2019.
- [65] R. Navigli and S. P. Ponzetto, "Babelnet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network," *Artificial intelligence*, vol. 193, pp. 217–250, 2012.
- [66] N. Mehrabi, P. Zhou, F. Morstatter, J. Pujara, X. Ren, and A. Galstyan, "Lawyers are dishonest? quantifying representational harms in commonsense knowledge resources," in *EMNLP*, 2021.
- [67] J. L. Fleiss and J. Cohen, "The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability," *Educational and psychological measurement*, 1973.
- [68] B. J. Devereux, L. K. Tyler, J. Geertzen, and B. Randall, "The centre for speech, language and the brain (CSLB) concept property norms," *Behavior research methods*, 2014.
- [69] K. Guu, K. Lee, Z. Tung, P. Pasupat, and M.-W. Chang, "Realm: Retrieval-augmented language model pre-training," in *ICML*, 2020.
- [70] F. Petroni, P. Lewis, A. Piktus, T. Rocktäschel, Y. Wu, A. H. Miller, and S. Riedel, "How context affects language models' factual predictions," in *AKBC*, 2020.
- [71] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "Albert: A lite Bert for self-supervised learning of language representations," in *ICLR*, 2020.