

Imitation of Life: A Search Engine for Biologically Inspired Design

Hen Emuna¹, Nadav Borenstein², Xin Qian³, Hyeonsu Kang⁴,
Joel Chan³, Aniket Kittur⁴, Dafna Shahaf¹

¹The Hebrew University of Jerusalem

²University of Copenhagen

³University of Maryland

⁴Carnegie Mellon University

hen.emuna@mail.huji.ac.il, nadav.borenstein@di.ku.dk, xinq@umd.edu, hyeonsuk@cs.cmu.edu, joelchan@umd.edu, nkittur@cs.cmu.edu, dshahaf@cs.huji.ac.il

Abstract

Biologically Inspired Design (BID), or Biomimicry, is a problem-solving methodology that applies analogies from nature to solve engineering challenges. For example, Speedo engineers designed swimsuits based on shark skin. Finding relevant biological solutions for real-world problems poses significant challenges, both due to the limited biological knowledge engineers and designers typically possess and to the limited BID resources. Existing BID datasets are hand-curated and small, and scaling them up requires costly human annotations.

In this paper, we introduce BARCODE (Biological Analogy Retriever), a search engine for automatically mining bio-inspirations from the web at scale. Using advances in natural language understanding and data programming, BARCODE identifies potential inspirations for engineering challenges. Our experiments demonstrate that BARCODE can retrieve inspirations that are valuable to engineers and designers tackling real-world problems, as well as recover famous historical BID examples. We release data and code; we view BARCODE as a step towards addressing the challenges that have historically hindered the practical application of BID to engineering innovation.

1 Introduction

Nature is a rich source of inspiration for creative problem-solving. For example, engineers addressing the bullet train’s sonic boom problem drew inspiration from kingfishers, who dive without splashing water (Linic et al. 2021). By re-designing the train’s front after a kingfisher’s beak, they reduced booms and improved speed while conserving energy.

Biologically Inspired Design (BID), or biomimicry, is a design approach that derives solutions by drawing analogies to nature’s strategies and systems (Fu et al. 2014; Lurie-Luke 2014; Dumanli and Savin 2016; Pawlyn 2019; San Ha and Lu 2020).

In the past two decades, there has been a growing interest in BID (Verbrugghe, Rubinacci, and Khan 2023). However, despite its potential, it is still not widely applied. One reason might be that finding relevant biological solutions is time-consuming due to limited biological knowledge among designers and engineers (Vattam and Goel 2011a). Analogies

from biology often rely on serendipity (Vattam and Goel 2013) – in the bullet train case, one of the engineers was an amateur birdwatcher (Kobayashi 2005). Moreover, while biological knowledge is abundant, bio-inspirational information is scarce. An audit of potential inspirations, involving heuristically filtered scientific articles from PubMed, revealed that only 3.36% (from a sample of 3416) provided insights for design problems (Kaiser et al. 2012).

Several hand-curated databases of BID exist. AskNature (asknature.org 2018) is an online dataset of biological phenomena, categorized by functional keywords. Design by Analogy to Nature Engine (DANE) provides a library of biological descriptions represented in the Structure-Behavior-Function (SBF) framework (Vattam et al. 2011; Goel, Rugaber, and Vattam 2009). However, the hand-curation of these databases is an extremely time-consuming and costly process; it takes 40–100 hours to describe a single biological organism in their framework. Thus, there is a need for automatic, *scalable* methods to expand the coverage of these datasets (Vandevonne et al. 2016).

More recently, Zhao et al. demonstrated how crowdsourcing annotations of scientific articles containing biomimicry-related content can facilitate a supervised learning approach to classifying scientific articles (Zhao et al. 2018; Vattam and Goel 2011b). However, generating labels for a large dataset can again be expensive and difficult.

Another work relies on hard-coded rules to seek bio-inspiration (Cheong and Shu 2014). However, the set of rules is small and limited. Furthermore, they match only the identified causal-function verbs, which are often too broad. For instance, a query like “reduce glare” only matches on “reduce”, possibly leading to irrelevant matches that do not refer to “glare”.

In this work, we introduce BARCODE (Biological Analogy Retriever), a *scalable* and *automatic* end-to-end system for mining bio-inspiration from unstructured data written in natural language. A recent work in analogy mining has shown how mining functional aspects such as the purposes and mechanisms of products can support the search and discovery of products that match on an abstract level (Hope et al. 2022). We incorporate recent advances in natural language understanding and data programming to extract functional aspects relevant to our settings. We show our system **retrieves valuable bio-inspirations** for human de-

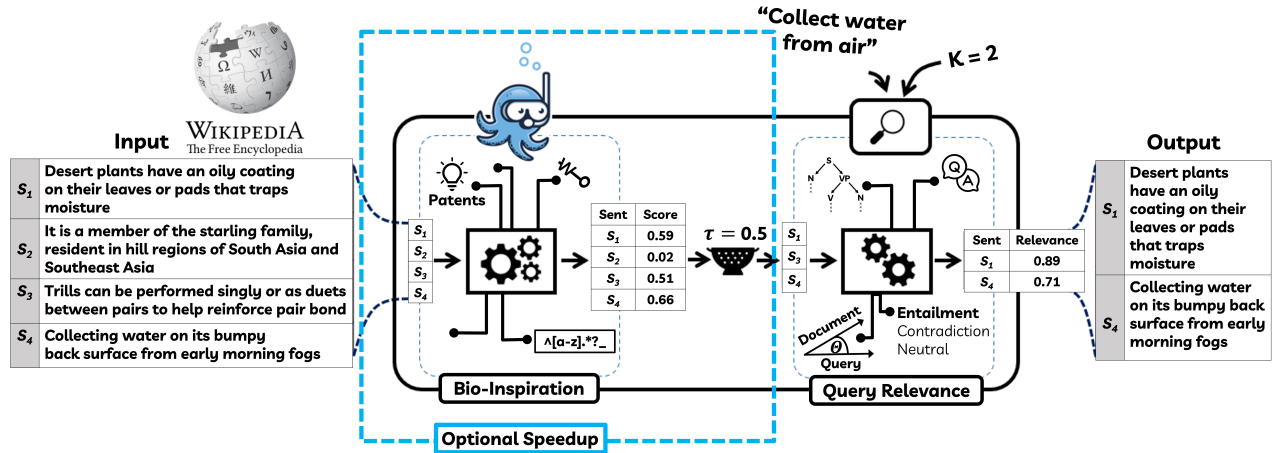


Figure 1: An overview of BARCODE. The bio-inspiration module, powered by Snorkel, processes a large corpus of biological documents, like Wikipedia, and generates a bio-inspiration score for each sentence in the corpus. The score measures the likelihood that the sentence describes a biological solution to a problem and is generated using linguistic labeling functions and external data sources like patent database. The query-relevance module then filters low-scoring sentences and, given a query, returns a ranked list of length k of relevant candidate sentences sorted by their relevance to the query.

sign problems, both in a **retrieval experiment** and with **real experts** looking for inspiration to solve their own research problems. We also demonstrate how BARCODE can **recover famous historical biomimicry examples**. We release data and code at <https://github.com/emunatool/BARcode-BioInspired-Search>.

We hope BARCODE will help to scale existing bio-inspiration datasets (e.g., AskNature) and aid designers and engineers in addressing real-world problems.

2 Problem Formulation

Given a corpus of documents \mathcal{D} written in a natural language and a query q describing a challenge (e.g., “collect water from humid air”), our goal is to return a ranked list of sentences from the corpus containing possible sources of inspiration. Specifically, we are looking for sentences describing an organism facing a similar *challenge*, and potentially the *strategy* it uses to overcome it. For example, the sentence “Desert plants have an oily coating on their leaves or pads that traps moisture, thereby reducing water loss” includes a challenge (trapping moisture/reducing water loss) and hints at a strategy (oily coating).

In our formulation, we are inspired by a recent work in analogy mining that demonstrated how mining functional aspects (such as the purposes and mechanisms of products) can support the search for analogous products that match on an abstract level (Hope et al. 2022). We adapt this idea into our setting, focusing on the aspects of challenge and strategy.

BARCODE can work with any corpus containing biological knowledge. For this paper, we focus on *Wikipedia* because of its accessibility and broad coverage of diverse biological topics (Wikimedia.org 2019). The corpus contains 780,949 sentences, collected from 27,640 articles listed un-

der the category “Articles with ‘species’ microformats”¹.

3 Mining Biological Analogies

We start by describing our algorithm (Section 3.1), and then discuss speedups in Section 3.2. An overview of our system can be seen in Figure 1. Our algorithm consists of two phases: First, an optional speedup phase to identify sentences that are likely to contain a challenge an organism faces (and potentially a coping strategy). Next, we **match and rank** the most relevant sentences to q .

3.1 Calculating Relevance Score

Here, our goal is to identify the most relevant sentences that would provide useful bio-inspiration for addressing the user’s challenge. Our system is designed to handle short queries describing a specific challenge, e.g., “collect water from humid air”. As sentences from the corpus can be long and convoluted, we start by extracting candidate phrases from the sentences in the form of verb-object pairs (e.g., “trap moisture”).

Extracting Candidate Phrases To extract verb-object pairs, we use two complementary techniques. The first technique employs Semantic Role Labeling (SRL) (Gildea and Jurafsky 2002), specifically we use QA-SRL (FitzGerald et al. 2018). QA-SRL is a model that takes a sentence as input and generates question-answer pairs related to the verbs². For example, applying QA-SRL on the sentence from Section 2: Q: ‘What does something trap?’,

¹We used spaCy’s NLP engine (Honnibal and Montani 2017) to process the articles.

²We omit ‘When’ and ‘Who’ questions and those not ending with a verb. Based on experimentation with the QA-SRL model, these question types usually don’t align with challenge-related verb-object pairs.

Organism	Sentence	Extraction Method
prevent sinking		
Ctenophora	If they enter less dense brackish water, the ciliary rosettes in the body cavity may pump this into the mesoglea to increase its bulk and decrease its density, to avoid sinking	QA-SRL
Barychelidae	Others can avoid drowning by trapping air bubbles within the hairs covering their bodies	QA-SRL
Pelican	The air sacs serve to keep the pelican remarkably buoyant in the water and may also cushion the impact of the pelican’s body on the water surface when they dive from flight into water to catch fish	DEP
Cephalopod	Other cephalopods use ammonium in a similar way, storing the ions as ammonium chloride to reduce their overall density and increase buoyancy	DEP
collect water from humid air		
Stenocara Gracilipes Yucca	Facing into the breeze, with its body angled at 45 degrees, the beetle catches fog droplets on its hardened wings	Both
	Some desert plants have an oily coating on their leaves or pads that traps moisture , thereby reducing water loss	QA-SRL
Kangaroo rat	To reduce loss of moisture through respiration when sleeping, a kangaroo rat buries its nose in its fur to accumulate a small pocket of moist air	QA-SRL

Table 1: Selected sentences from the Top 15 retrieved by BARCODE for two queries. The phrases that are used to match to the query are highlighted in bold.

A: ‘moisture’. We then convert each question-answer pair into the format ‘[verb] [answer]’³ (e.g., ‘trap moisture’).

However, QA-SRL has limitations. It fails to extract certain challenge-related phrases, especially those with auxiliary verbs (e.g., “keep buoyant”); it tends to produce verbose phrases (“store ions to reduce overall density and increase buoyancy”). To address these limitations, we employ a second technique, pattern-matching on dependency trees with spaCy (Honnibal and Montani 2017). Dependency parsing offers greater flexibility in extracting verb-object pairs, allowing connections between non-adjacent verbs and objects. Through experimentation, we have found 10 patterns (see Appendix A.1⁴). By employing these two techniques, we extracted 2,829,204 candidate phrases from sentences.

Matching and Ranking Next, we employ a two-step approach to assess the similarity between these candidate phrases and the query. We employed SBERT, designed for semantic search (Reimers and Gurevych 2019)⁵, and DeBERTa (He et al. 2020)⁶ as the inference model.

Sentence embedding models, like SBERT, effectively capture semantic meaning, but they have trouble handling antonyms (e.g., “increase” and “decrease” might be considered similar) (Wen-tau Yih and Platt 2012; Silke Scheible and Springorum 2013), which is problematic for our use case. Additionally, these models struggle with causality and

common sense, causing them to overlook relevant matches (e.g., make the connection between “stay moist” and “reduce water loss”).

To address these limitations, we add to our pipeline a Natural language inference (NLI) model (I.Dagan and B.Magnini 2006). NLI is the task of determining whether a hypothesis is true (entailment), false (contradiction), or undetermined (neutral) given a premise, and can address the shortcoming of the embedding similarity techniques (Vah-tola, Creutz, and Tiedemann 2022; LoBue and Yates 2011).

We chose DeBERTa, which has shown good performance on prominent NLI benchmarks. Candidate phrases form the premise and the query forms the hypothesis, resulting in three scores for every pair (entailment, neutral, and contradiction). We note that DeBERTa indeed manages to identify that “stay moist” implies “reduce water loss” and that “increase water loss” contradicts “reduce water loss”.

Using the inference model is computationally costly. Thus, we restrict ourselves to the top 4,000 phrases according to semantic similarity (SBERT). Looking at the data, 4000 seems to provide a large safety margin, so we are not likely to filter out any promising candidate.

To create a single weighted score, we train a simple classifier to classify pairs as *relevant* or *irrelevant* based on the NLI scores and SBERT score (See Appendix A.1 for details).

We retrieve the top sentences according to our classifier. In Table 1, we provide examples of retrieved sentences for two queries: “prevent sinking”, and “collect water from humid air”.

³[verb] represents the lemma form of the main verb in the question, and [answer] represents the span of the answer.

⁴For all appendix references, see the arXiv version of our paper: <https://doi.org/10.48550/arXiv.2312.12681>

⁵<https://huggingface.co/sentence-transformers/multi-qa-mpnet-base-dot-v1>

⁶<https://huggingface.co/cross-encoder/nli-deberta-v3-base>

3.2 Speeding Up the Algorithm

We built BARCODE to work with any corpus containing biological knowledge. However, sentences that could serve as biomimetic inspiration are typically very rare in such corpora. When the corpus is very large, running the algorithm could become computationally expensive. Thus, in this section, we suggest filtering heuristics to narrow down the sentences we consider as candidates to match with a query.

To implement the filtering algorithm, we leverage recent advancements in **data-programming**. Data-programming relies on three key elements: 1) extracting *candidates* for labeling, 2) generating *labeling functions* that express domain heuristics (distant supervision), and 3) using a *generative model* to handle noise and correlations among labeling functions, resulting in probabilistic labels for all candidates.

We use Snorkel (Ratner et al. 2017), defining a candidate as three text spans (*Strategy*, *Solver*, and *Problem*). Note that sentences might include several candidates.

We employed ten clausal patterns to detect candidates. Seven patterns were adapted from Cheong and Shu (2014) (e.g., “the *Strategy* serves as a clausal subject for the *Problem* and is in gerund form”). After experimenting with the data, we added three more patterns: (1) *Problem* functions as an adverbial clause modifier, modifying the *Strategy*, (2) *Problem* is the main clause head of a relative clause, in which the *Strategy* appears in the main clause, and (3) *Problem* is the main clause head of an adjectival clause, where the *Strategy* serves as the sentence’s root. Out of 494,640 dataset entries, approximately 234,000 contained one or more clausal patterns and were used as candidates.

Defining Labeling Functions (LF). We defined labeling functions (LFs) to assign each candidate to one of three labels: contains bio-inspiration, does not contain bio-inspiration, or unknown.

Some LFs look for cues that a sentence contains bio-inspiration:

- **Keywords:** The sentence contains the word “adaptation” or one of its conjugations. We notice that almost all the sentences that contain this word indeed describe an interesting analogy. For example, “*The eyes appear to be narrowly open due to the lowered upper eyelid, probably an adaptation to shield the eyes from the sun’s glare*”.
- **Distant Supervision – Patents:** The sentence contains a description of a known problem. Candidates that contain a known problem are also likely to contain a solution to this problem. We used a dataset of **patents** to construct a list of known problems (using what one might call “cross-corpus distant supervision”): First, we downloaded 20 million sentences from the “Claims” section of a random set of patents. Claims describe what the patent is supposed to do. Then, we extracted from this section strings in the form of “for [verb]-ing [noun]” using regular expressions and POS (Part Of Speech) tags. Coming from the claims section, those strings are likely to address the problem that the patent was designed to solve. For example, “*A surface cleaning apparatus ... comprising a second collecting apparatus ... for collecting liquid from the surface*”.

We listed the 2,000 most common verb-noun pairs and used them as our known problems list. This list contains pairs like “tilt movement”, “extract signal” and “sense light”. We say that a candidate describes a known problem if it contains such a noun-verb pair, which are close to one another in the sentence. Note that many of the problems coming from patents are unlikely to be solved by animals (e.g., “encode information”).

- **Auxiliary Verbs:** The sentence contains an auxiliary verb like “allow”, “help” or “enable”. Those verbs exist in many relevant candidates that describe a biological trait, that is used by an organism to solve a problem. The problem, in this case, will appear right after the auxiliary verb, as seen in the following sentence: *The webbing between the toes increases the area of the foot and helps propel the frog powerfully through the water*”.

Other LFs look for cues that a sentence is *unrelated* to biological strategies:

- **Verbs with No Biology Relation:** The main verb of the sentence is not biology-related. Among those, we can find verbs such as “dance”, “read”, “explain”, and “discover”. This list of verbs was hand-crafted from the 200 most common verbs in the English language and contains 36 entries.
- **Unlikely Entities:** The sentence contains a component that is not expected to be found in functional sentences. For example, if the sentence contains the name of a person or organization, a date, a pronoun, or a symbol (such as &, @, or \$), it is unlikely that it describes a biological solution to a problem.

Next, Snorkel predicts the likelihood that a candidate is bio-inspirational. Note that this process is performed as a pre-processing step and does not depend on the query.

See Table 2 for some of Snorkel’s scores. We set the filtering threshold to $\tau = 0.5$ (after an examination of the data, this seems to nicely balance precision and recall), leaving us with 23,553 sentences (3% of the total data). In the evaluation section we show that running on this much-smaller set of sentences does not harm performance much.

4 Evaluation

In this section, we report our evaluation of the performance of BARCODE’s bio-inspiration score module, as well as an end-to-end evaluation of the system’s ability to retrieve biomimetic inspirations.

Our research questions are as follows:

- **RQ1:** Can our approach provide inspiration from nature to solve human design problems?
- **RQ2:** Is our algorithm robust to query variations?
- **RQ3:** Does the bio-inspiration module predict whether a sentence contains a challenge?

A Note on Large Language Models. Large language models have recently gained immense popularity, achieving state-of-the-art results across a variety of downstream tasks. We have initially experimented with applying ChatGPT (<https://openai.com/chatgpt>) to our problem, but quickly

Candidate sentence	Snorkel's Score
<i>Peregrine falcon</i> — The air pressure from such a dive could possibly damage a bird's lungs, but small bony tubercles on a falcon's nostrils guide the powerful airflow away from the nostrils, enabling the bird to breathe more easily while diving by reducing the change in air pressure	0.935
<i>Isopoda</i> — The dorsal (upper) surface of the animal is covered by a series of overlapping, articulated plates which give protection while also providing flexibility	0.802
<i>Yucca</i> — Some desert plants have an oily coating on their leaves or pads that traps moisture , thereby reducing water loss	0.598
<i>Pigeon guillemot</i> — Trills can be performed singly or as duets between pairs; if performed as a duet then the call also functions to help reinforce pair bond	0.469
<i>Morgan horse</i> — By the 1870s, however, longer-legged horses came into fashion, and Morgan horses were crossed with those of other breeds	0.213
<i>Common hill myna</i> — It is a member of the starling family (Sturnidae), resident in hill regions of South Asia and Southeast Asia	0.0002

Table 2: Examples of Snorkel's score for the filtering of the dataset. *Challenge* and *strategy* are highlighted in bold.

found out that it 1) tends to hallucinate facts that we could not corroborate (“Termites build mounds that capture and collect moisture from the air through a process known as passive condensation”), making it hard to trust, and 2) has limited coverage (and when pressed to come up with more organisms, simply repeats the first ones). Thus, we decided that an information-retrieval setting that covers the entire Wikipedia and can point our users to the source of each claim is better suited for our use case.

4.1 Evaluating Retrieval Quality

In this section, we evaluate the retrieval quality of our system to tackle (RQ1) — can our algorithm provide relevant inspirations to solve human design problems?

Reality Check: Famous Historical Examples To assess our system's retrieval quality, we first try it on famous historical biomimicry examples. We scoured the 'INNOVATIONS' section of AskNature, which outlines nature-inspired breakthroughs, and sampled five examples to formulate queries and search for them. The results were:

- Capturing cooling tower plumes to reduce water usage, inspired by Namibian desert beetles. Ranked 8 for the query “collect water from humid air”.
- Residue-free adhesion technology, inspired by gecko

feet. Ranked 9 for the query “provide adhesion”.

- Speedo's swimsuit for faster swimming, inspired by shark skin. Ranked 10 for the query “reduce fluid drag”.
- Passively cooled buildings, inspired by termite mounds. Ranked 14 for the query “provide self-regulating ventilation system”.
- Train's front inspired by a kingfisher's beak, not found. We checked the kingfisher's Wikipedia page, and it does not mention the beak's hydrodynamic structure.

We note that famous historical examples might be easier to find, as Wikipedia is more likely to contain a sentence about them; thus, we consider this a reality check and move on to further evaluation of the retrieval quality.

Retrieval Experiment: Collecting Queries To collect relevant innovation challenges, we used titles from AskNature's “Biological Strategies” section. These titles outline how organisms address challenges (e.g., “Adaptive camouflage helps blend into the environment”). We converted these titles into query format (verb-object pairs, e.g., “blend into environment”). This resulted in 24 queries.

To assess the robustness of our system (RQ2) (Section 4.2), we generated more technical variations for 18 queries (e.g., “repel water”/“prevent water absorption”).

Baseline. As we could not access the code for the work of Cheong and Shu (2014), we compared our system with a search engine (which is the most likely tool to be used by people looking for biomimetic inspiration today). We used Elasticsearch (v7.13.2), a widely used search engine built on top of the Lucene library. We used Elasticsearch's default relevance scoring function. Both for our algorithm and for the baseline, we consider two settings:

- **Entire data:** Used all sentences in the corpus.
- **Filtered data:** Only sentences passing the threshold $\tau = 0.5$, corresponding to aggressive pruning of the data (keeping 3% of the entire data). See Section 3.2.

Experiment Design. We collected the top 15 results per query from BARCODE and the baseline for both settings, resulting in 2074 unique sentences. We used crowdsourcing to determine whether each sentence contains a *challenge* and/or *strategy*. Annotators were presented with an innovation challenge (query) along with a sentence about an organism (see Appendix A.3). Annotators were asked to assess whether the sentence indicates: (a) the organism faces a similar challenge, (b) provides a strategy to tackle the challenge, (c) both, or (d) is irrelevant.

We hired 29 Amazon Mechanical Turk (AMT) crowd workers who passed a qualification test. They received \$0.11 per task. Each sentence was annotated by 3 workers, and the presence of a challenge and/or strategy was determined by a majority vote.

The annotators showed moderate agreement, as indicated by Fleiss' Kappa scores of 0.42 for challenge and 0.55 for strategy evaluation. This might indicate the difficulty of the task: sentences sometimes lack context, and challenges and strategies are sometimes implicit. We note that some annotators preferred to mark only the more explicit sentences, while others did not show such preference.

Results. Table 3 compares BARCODE with the baseline,

both on the entire data and filtered data, using two information retrieval metrics: *Precision @k* and *Normalized Discounted Cumulative Gain (NDCG)@k* at two cut-offs ($k=7$, $k=15$). For NDCG, we use binary relevance.

Our system outperforms the baseline in all metrics. Most importantly, on the entire data, 72% of the top 7 sentences and 69% of the top 15 sentences retrieved by BARCODE contain a strategy that could be inspirational for solving the challenge (which is quite high in absolute terms). In comparison, the baseline only yields 36% for the top 7 and 29% for the top 15.

The situation is similar for sentences that contain a challenge, with our system outperforming the baseline. We note that the absolute percentage of those sentences is lower (which surprised us). Looking at the data, this seems to stem from the fact that annotators tended not to mark sentences where the challenge was somewhat implicit.

On the filtered data, our system again beats the baseline on all metrics. There is a trade-off between speed (filtered data is only 3% of the entire data) and performance. For both BARCODE and the baseline, filtering does not hurt performance much for the top 7 but starts to be more pronounced at the top 15. We believe that a less aggressive threshold would have narrowed the gap; we leave the problems of finding automated ways to pick a threshold for future work.

We used the Mann-Whitney U test to assess the statistical significance of differences in precision between BARCODE and baseline on the entire data, and also between BARCODE on the entire data and the filtered data. This analysis was conducted separately for *challenge* and *strategy* at two cut-off points: $k=7$ and $k=15$, covering all queries (with α of 0.05 before Bonferroni correction). Comparing BARCODE to the baseline, all tests were statistically significant (p-values: $2.26e-7$ for $k=7$, $2e-10$ for $k=15$ for strategy, and $5.96e-6$ for $k=7$, $2.39e-8$ for $k=15$ for challenge). Differences between BARCODE on the entire data and filtered data were not significant, except for strategy at $k=15$ (p-value= $3.88e-3$). This supports our observations about the utility of filtering.

Next, we take a look at query paraphrases. Figure 2 shows Precision@k results for AskNature queries and their variations. Notably, the gap between BARCODE and the baseline is greater for the paraphrases. We hypothesize that this happens because the paraphrases use more technical terms than the original ones, and also sometimes require commonsense to understand the connection, which might be captured by our NLI component (but not by a regular search engine).

Error Analysis. We examined the 10 queries achieving the worst performance, yielding 106 sentences annotated as irrelevant. Among these, 29% were due to annotators overlooking challenge/strategy within sentences, often when it is implicit. In 39% of cases, the matched candidate phrase did not capture the intended challenge, like “detect carbon dioxide” being matched to the query “capture carbon dioxide”. In 27% of cases, the matched candidate phrase was not actually about the organism. In 5% of the cases, the sentences lack bio-inspirational content.

To conclude, BARCODE **outperforms the baseline** across the board, and can **provide useful inspirations** from

	Method	P	NDCG
Challenge	Baseline Filtered @7	0.272	0.554
	Baseline Entire @7	0.282	0.536
	BARCODE Filtered @7	0.541	0.736
	BARCODE Entire @7	0.565	0.787
	Baseline Filtered @15	0.224	0.557
	Baseline Entire @15	0.24	0.554
	BARCODE Filtered @15	0.435	0.745
	BARCODE Entire @15	0.541	0.798
Strategy	Baseline Filtered @7	0.286	0.565
	Baseline Entire @7	0.357	0.647
	BARCODE Filtered @7	0.568	0.763
	BARCODE Entire @7	0.718	0.884
	Baseline Filtered @15	0.252	0.579
	Baseline Entire @15	0.295	0.658
	BARCODE Filtered @15	0.514	0.772
	BARCODE Entire @15	0.694	0.877

Table 3: Precision (P)@k and NDCG@k for all queries. Both BARCODE outperforms the baseline in every metric, both on the entire data and on filtered data. Filtering does not hurt performance much for the top 7, but starts to be more pronounced at the top 15.

nature to a large variety of queries (see also experiments with experts in Section 5). Our filtering heuristics are able to drastically reduce the size of the data before performance starts to decline.

4.2 Robustness to Query Variations

We now address **RQ2**. To test the robustness of our system to query variations, we created alternative phrasings for 18 queries from AskNature, resulting in 18 pairs of queries. We evaluated these pairs using two metrics: First, we looked at the number of shared results among the top 15 retrieved sentences for each pair. Second, we used the Rank-biased Overlap (RBO) (Webber, Moffat, and Zobel 2010), which compares two ranked lists, taking into account both the shared items and their order, providing a score between 0 and 1.

Results. On the full dataset, BARCODE averages an RBO of 0.13 and of 2.83 shared items. This can be attributed to the size of the dataset it operates on and the number of possible good answers. On the filtered data, the average is RBO of 0.31 and 6.8 shared items (hinting at similar items ranked at different places). In other words, the system is more robust on the filtered data.

4.3 Evaluating Bio-Inspiration Scores

We now aim to assess if the filtering module (Section 3.2) identifies the kind of sentences we are after (**RQ3**). We compare to Cheong’s rule-based approach (Cheong and Shu 2014), the previous state-of-the-art.

We lack a database of gold-standard labeled candidates, and the majority of sentences in our corpus are unlikely to contain challenge-strategy pairs. Thus, we randomly selected 4 verbs from the titles of AskNature articles. We ex-

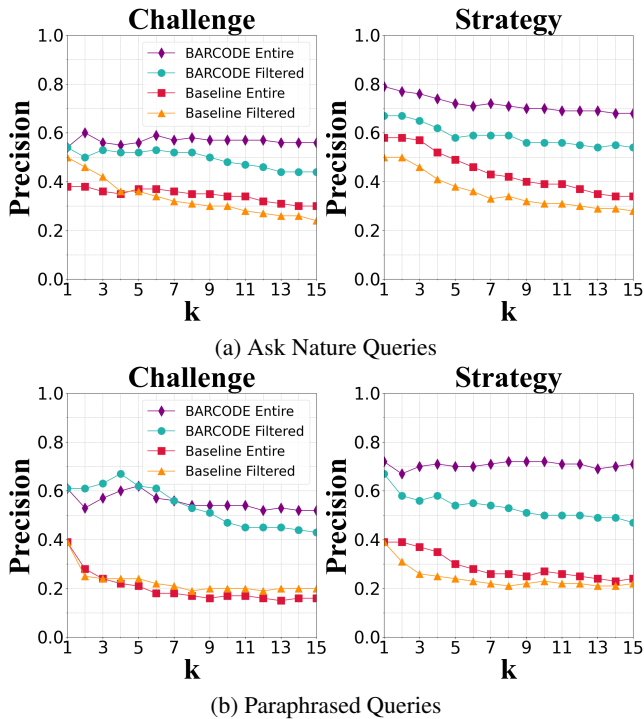


Figure 2: Precision @k across all queries. *Top*: AskNature queries. *Bottom*: Paraphrased queries. *Left*: sentences containing challenges similar to the query. *Right*: sentences containing strategies to tackle the challenge in the query. BARCODE outperforms the baseline both on the entire data and filtered data. The gap in performance is wider on the paraphrased queries, perhaps due to our NLI model.

Top K	Calculation type	BARCODE	Cheong
High	Strict	0.56	0.22
	Liberal	0.80	0.25
Low	Strict	0.13	0.15
	Liberal	0.28	0.33

Table 4: Comparing the performance of BARCODE and Cheong’s rule-based approach. Our approach outperforms Cheong’s method in top 10 score precision for both strict and liberal cutoffs. Recall is similar, with our approach having slightly fewer inspirational sentences in the bottom 10.

tracted all sentences containing those verbs from Wikipedia and computed their ranking according to Cheong and according to our method. We sampled ten sentences per verb from the top and the bottom of each ranked list, yielding a total of 160 sentences for evaluation. We evaluate the bottom of the lists as well to provide a rough estimate of recall.

Four members of the research team, blind to condition, scored the candidates on a scale of 0 (definitely does not contain a strategy to address a challenge) to 2 (definitely contains a strategy to address a challenge).

Inter-rater reliability was acceptable (average pairwise

	Queries	#	Known	Interest
E1	maintain aqueous temperature	30	4	2 (1)
	regulate temperature	15	0	1 (0)
	provide insulation in extreme temperatures	15	0	1 (0)
E2	conserve heat	15	0	1 (0)
	measure metabolic activity	25	1	6 (1)
	sense lactate	15	0	3 (2)
E3	assemble cells	15	2	3 (2)
	sense electrical fields	25	0	10 (3)
	recognize electrical activity in a specific radius	15	0	2 (2)

Table 5: Ideation experiment. For each expert query, we show the number of results, number of results they were already familiar with (and were valuable), number of results they found somewhat interesting or interesting (score 1 or 2). In parenthesis, the number of very interesting results (score 2). 33% of the relevant results (known or interesting) were considered highly interesting, and 21% of them were already known to the experts.

Spearman’s $\rho = .60$). We computed a score for each match by averaging the 4 ratings, at both strict (average = 2) and liberal (average ≥ 1) cutoffs.

Results. Table 4 demonstrates that our approach substantially outperforms Cheong’s rule-based approach in terms of precision for both strict (0.56 vs. 0.22) and liberal (0.80 vs. 0.25) cutoffs in the top 10. Our approach exhibits marginally fewer inspirations in the bottom 10 compared to the rule-based method, hinting that it does not compromise recall.

5 Ideation Experiment: Case Studies

To evaluate whether our results extend to real-world scenarios, we recruited three experts engaged in innovative bio-engineering projects: E1 sought inspiration for maintaining a stable temperature of liquid inside a tube, E2 wished to measure metabolic activity and cell assembly, and E3 sought new ways to record electrical activity from cell groups.

We guided the experts to formulate queries in a “verb-object” format, and instructed them to abstract their needs (rather than use extremely specific technical terms from their domain). We used the sped-up version of our algorithm to allow live interaction with the search engine.

Known Directions. The experts were pleasantly surprised by the search results BARCODE returned and noted several that matched their own findings. : E1 mentioned, “...this would be the best solution for us... [wrapping their tube with another tube and maintaining a high level of moisture between them] ...we already thought about this idea, it is similar to the process used to pasteurize milk and heat blood samples”. Similarly, E2 stated, “It is something that I already know and try to apply...”.

Overall, 21% of the relevant results were familiar to the

experts (see Table 5). We take this as a positive signal, indicating the BARCODE does indeed manage to identify directions that experts had independently considered.

Novel and Inspirational. BARCODE inspired experts to generate diverse ideas, revealing new perspectives and insights. For instance, E2 expressed enthusiasm, stating, “this is the best result so far...it (catfish) has a system through its facial tasting sensors can detect metabolites... and I am looking for sensors for metabolites... I am intrigued to know how this mechanism works... also this mechanism seems to be constantly active, unlike current techniques”. Regarding the query of “measure metabolic activity”, E2 stated “...this is very interesting... I did not expect this result... it mentions that a low metabolic rate reduces mortality and increases longevity... I would be interested in exploring this”. In another instance, E1 dwelled on a moth’s insulation layer used to prevent overheating. E1 highlighted not having considered such a structure before, and that it could serve as a more effective insulation layer. Experts rated each result on a scale from 0 (not interesting) to 2 (very interesting).

Table 5 shows the search queries, number of BARCODE results read by the expert during the session, and number of entries found interesting by experts (in parenthesis: number of very interesting results). We note that all but three of the queries led to at least one inspiration deemed “very interesting” by the expert.

After the interview, we thanked the experts for their participation. No other compensation was provided.

6 Conclusions and Future Work

Biologically Inspired Design (BID), or Biomimicry, is a problem-solving methodology that applies analogies from nature to solve challenges. Finding relevant biological solutions for real-world problems poses significant challenges, both due to the limited biological knowledge engineers and designers typically possess and the scarcity of sources of inspiration. Current BID datasets are manually curated, and scaling them up requires expensive manual effort.

This paper introduces BARCODE (Biological Analogies Retriever), a flexible search engine that automatically mines bio-inspirations from the web. Leveraging advancements in natural language understanding and data programming, BARCODE identifies potential solutions rooted in nature for various engineering challenges. Our experiments confirm BARCODE’s ability to extract valuable inspiration from large corpora, including famous historical examples of biomimicry. BARCODE was able to find promising inspirations for experts in engineering for their research problems.

We note that across our piloting of the system, we noticed substantial variance in performance across queries. Some of this might be due to the idiosyncrasies of the composition of our corpus; however, our experiences with the case study experts, in particular, suggest that exploring interactive mechanisms for query reformulation and expansion (perhaps with the help of large language models) might significantly improve the user experience.

In the future, we hope to apply BARCODE to different, more detailed sources of biological data, such as scientific

papers. In this work, we have chosen to concentrate on retrieving inspirations; an exciting direction for future exploration is integrating BARCODE with generative artificial intelligence technologies to generate a sketch of the solution, adapting the biological strategy retrieved by BARCODE to the user’s problem. We see BARCODE as a step towards addressing the challenges that have historically hindered the practical application of BID to engineering innovation.

Acknowledgements

We would like to thank Hyadata Lab and NLP-HUJI members for their thoughtful remarks, and all the participants in our experiments for their efforts. We also thank the reviewers for their insightful comments. This work was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant no. 852686, SIAM), Office of Naval Research, and Carnegie Mellon Center for Knowledge Acceleration. Hen Emuna was supported by Sam and Betty Sandler and the Samuel and Lottie Rudin Scholarship.

References

- asknature.org. 2018. AskNature - Innovation inspired by nature. <http://asknature.org>. Accessed: 2018-09-09.
- Cheong, H.; and Shu, L. H. 2014. Retrieving Causally Related Functions From Natural-Language Text for Biomimetic Design. *Journal of Mechanical Design*, 136(8): 081008–081008–10.
- Dumanli, A. G.; and Savin, T. 2016. Recent advances in the biomimicry of structural colours. *Chemical Society Reviews*, 45(24): 6698–6724.
- FitzGerald, N.; Michael, J.; He, L.; and Zettlemoyer, L. 2018. Large-Scale QA-SRL Parsing. *CoRR*, abs/1805.05377.
- Fu, K.; Moreno, D.; Yang, M.; and Wood, K. L. 2014. Bio-Inspired Design: An Overview Investigating Open Questions From the Broader Field of Design-by-Analogy. *Journal of Mechanical Design*, 136(11): 111102–111102–18.
- Gildea, D.; and Jurafsky, D. 2002. Automatic labeling of semantic roles. *Computational linguistics*, 28(3): 245–288.
- Goel, A. K.; Rugaber, S.; and Vattam, S. 2009. Structure, behavior, and function of complex systems: The structure, behavior, and function modeling language. *Ai Edam*, 23(1): 23–35.
- He, P.; Liu, X.; Gao, J.; and Chen, W. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Honnibal, M.; and Montani, I. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. *To appear*.
- Hope, T.; Tamari, R.; Hershcovich, D.; Kang, H. B.; Chan, J.; Kittur, A.; and Shahaf, D. 2022. Scaling Creative Inspiration with Fine-Grained Functional Aspects of Ideas. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–15.

- I.Dagan, O.; and B.Magnini. 2006. “The pascal recognising textual entailment challenge”. In *Machine Learning Challenges. Evaluating Predictive Uncertainty, Visual Object Classification, and Recognising Tectual Entailment*, 177–190.
- Kaiser, M.; Hashemi Farzaneh, H.; Lindemann, U.; et al. 2012. An approach to support searching for biomimetic solutions based on system characteristics and its environmental interactions. In *DS 70: Proceedings of DESIGN 2012, the 12th International Design Conference, Dubrovnik, Croatia*, 969–978.
- Kobayashi, K. 2005. JFS Biomimicry Interview Series: No.6 ‘Shinkansen Technology Learned from an Owl?’ - The story of Eiji Nakatsu. https://www.japanfs.org/en/news/archives/news_id027795.html. Accessed: 2018-09-09.
- Linic, S.; Lucanin, V.; Zivkovic, S.; Rakovic, M.; and Puharic, M. 2021. Experimental and numerical methods for concept design and flow transition prediction on the example of the bionic high-speed train. In *Experimental and Computational Investigations in Engineering: Proceedings of the International Conference of Experimental and Numerical Investigations and New Technologies, CNNTech 2020*, 65–82. Springer.
- LoBue, P.; and Yates, A. 2011. Types of Common-Sense Knowledge Needed for Recognizing Textual Entailment. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 329–334. Portland, Oregon, USA: Association for Computational Linguistics.
- Lurie-Luke, E. 2014. Product and technology innovation: What can biomimicry inspire? *Biotechnology advances*, 32(8): 1494–1505.
- Pawlyn, M. 2019. *Biomimicry in architecture*. Routledge.
- Ratner, A.; Bach, S. H.; Ehrenberg, H.; Fries, J.; Wu, S.; and Ré, C. 2017. Snorkel: Rapid Training Data Creation with Weak Supervision. *ArXiv e-prints*.
- Reimers, N.; and Gurevych, I. 2019. Sentencebert: Sentence embeddings using siamese bertnetworks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics*.
- San Ha, N.; and Lu, G. 2020. A review of recent research on bio-inspired structures and materials for energy absorption applications. *Composites Part B: Engineering*, 181: 107496.
- Silke Scheible, S. S. i. W.; and Springorum, S. 2013. Uncovering distributional differences between synonyms and antonyms in a word space model. In *In Proceedings of the 6th International Joint Conference on Natural Language Processing, Nagoya, Japan*, 489–497.
- Vahtola, T.; Creutz, M.; and Tiedemann, J. 2022. It Is Not Easy To Detect Paraphrases: Analysing Semantic Similarity With Antonyms and Negation Using the New SemAntoNeg Benchmark. In *Proceedings of the Fifth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, 249–262. Abu Dhabi, United Arab Emirates (Hybrid): Association for Computational Linguistics.
- Vandevenne, D.; Verhaegen, P.-A.; Dewulf, S.; and Duflou, J. R. 2016. SEABIRD: Scalable search for systematic biologically inspired design. *AI EDAM*, 30(1): 78–95.
- Vattam, S.; and Goel, A. 2013. Seeking bioinspiration online: a descriptive account. In *DS 75-7: Proceedings of the 19th International Conference on Engineering Design (ICED13), Design for Harmonies, Vol. 7: Human Behaviour in Design, Seoul, Korea, 19-22.08. 2013*.
- Vattam, S.; Wiltgen, B.; Helms, M.; Goel, A. K.; and Yen, J. 2011. DANE: Fostering Creativity in and through Biologically Inspired Design. In Taura, T.; and Nagai, Y., eds., *Design Creativity 2010*, 115–122. London: Springer London. ISBN 978-0-85729-224-7.
- Vattam, S. S.; and Goel, A. K. 2011a. Foraging for inspiration: understanding and supporting the online information seeking practices of biologically inspired designers. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 54860, 177–186.
- Vattam, S. S.; and Goel, A. K. 2011b. Semantically Annotating Research Articles for Interdisciplinary Design. In *Proceedings of the Sixth International Conference on Knowledge Capture, K-CAP ’11*, 165–166. New York, NY, USA: ACM. ISBN 978-1-4503-0396-5.
- Verbrugge, N.; Rubinacci, E.; and Khan, A. Z. 2023. Biomimicry in Architecture: A Review of Definitions, Case Studies, and Design Methods. *Biomimetics*, 8(1).
- Webber, W.; Moffat, A.; and Zobel, J. 2010. A Similarity Measure for Indefinite Rankings. *ACM Trans. Inf. Syst.*, 28(4).
- Wen-tau Yih, G. Z.; and Platt, J. C. 2012. Polarity inducing latent semantic analysis. In *In Proceedings of EMNLP*.
- Wikimedia.org. 2019. Wikipedia:Statistics. <https://en.wikipedia.org/wiki/Wikipedia:Statistics>. Accessed: 2019-01-30.
- Zhao, Y.; Baldini, I.; Sattigeri, P.; Padhi, I.; Lee, Y. K.; and Smith, E. 2018. Data Driven Techniques for Organizing Scientific Articles Relevant to Biomimicry. In *ACM/AAAI Artificial Intelligence, Ethics and Society (AIES) conference*.