BIOSPARK: An End-to-end Generative System for Biological-Analogical Inspirations and Ideation

Hyeonsu B. Kang[†], **David Chuan-en Lin**[†], **Nikolas Martelaro**[†], **Aniket Kittur**[†]
Yan-Ying Chen[‡], **Matthew K. Hong**[‡]

[†]Human-Computer Interaction Institute, Carnegie Mellon University, Pittsburgh, PA 15213
[‡]Toyota Research Institute, Los Altos, CA 94022

[†]`{hyeonsuk,chuanenl,nikmart,nkittur}@cs.cmu.edu` [‡]`{yan-ying.chen,matt.hong}@tri.global`

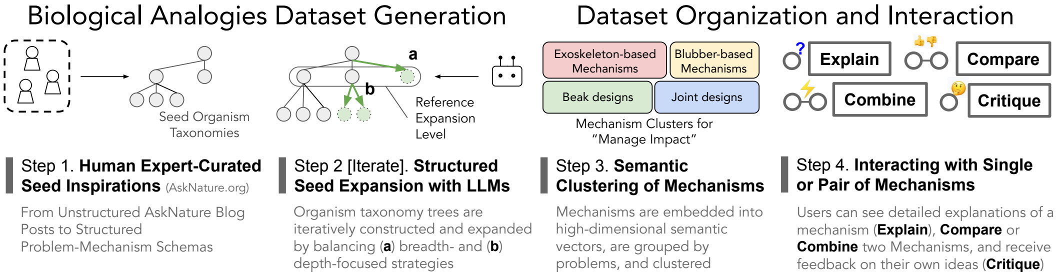
Biological Analogies Dataset Generation
Dataset Organization and Interaction
Seed Organism Taxonomies
a
b
Reference Expansion Level
Exoskeleton-based Mechanisms
Blubber-based Mechanisms
Beak designs
Joint designs
Mechanism Clusters for "Manage Impact"
?
Explain
Compare
Combine
Critique
Step 1. Human Expert-Curated Seed Inspirations (AskNature.org)
From Unstructured AskNature Blog Posts to Structured Problem-Mechanism Schemas
Step 2 [Iterate]. Structured Seed Expansion with LLMs
Organism taxonomy trees are iteratively constructed and expanded by balancing (a) breadth- and (b) depth-focused strategies
Step 3. Semantic Clustering of Mechanisms
Mechanisms are embedded into high-dimensional semantic vectors, are grouped by problems, and clustered
Step 4. Interacting with Single or Pair of Mechanisms
Users can see detailed explanations of a mechanism (Explain), Compare or Combine two Mechanisms, and receive feedback on their own ideas (Critique)

Figure 1: BIOSPARK consists of a *Dataset Generation Pipeline* and an *Organization and Interaction* interface. Dataset construction starts from a small set of seed mechanisms on AskNature (**Step 1**), expanded by iteratively constructing taxonomic hierarchies of organisms and diversifying using *breadth*- or *depth*-focused knowledge prompts (**Step 2**). Dataset is organized by semantic clusters (**Step 3**) and provided with interactive features that facilitate user learning and synthesis (**Step 4**).

1 Introduction

Nature provides a valuable source of inspirations [2, 26, 22] for novel design solutions to challenging engineering problems [23, 22]. A central cognitive mechanism underpinning this transfer of domains is *analogical processing* [8, 20, 12], which involves discovering and recognizing deep structural similarity between source and target analogs. To this end, recent work has shown that analogs with a similar purpose (*i.e.,* '*what problem it solves*'), but diverse mechanisms (*i.e.,* '*how it solves it*'), facilitate creative new ideas [17, 15].

Yet, the full potential of biological-analogical inspirations in engineering and design domains has proven difficult. Many bio-inspired designs are biased toward well-known organisms [3] and prior work has shown it can take on average three person-hours to sift through vast, fragmented information sources such as Google Scholar, Wikipedia, AskNature.org, or the open Web to find a single relevant article for a design problem [25, 24]. Human-expert-curated databases such as AskNature [7], DANE [13], SAPPhIRE [4], BioTRIZ [26], or rule-based databases (*e.g.,* [5]) can streamline searches, but the effort-intensive process of manual expert curation often results in data scarcity, limiting the diversity of organisms and mechanisms available for making analogies. Automatic generation of analogies with LLMs could create larger datasets for bio-inspired design, but they still have limited diversity on abstract concepts, even when tuning parameters to increase diversity such as logic suppression and increased sampling temperature [6]. Knowledge-augmented prompting provides an alternative design space to token-level manipulation, as explored in [1] for improving the factuality of Q&A tasks. We explore this design space for a different purpose of constructing a diverse biological mechanisms dataset, with a modular and structured knowledge augmentation approach.

Beyond sourcing biological analogies, it can be challenging for people to understand and recognize them when presented. People's cognitive biases towards surface similarity and working memory

37th Conference on Neural Information Processing Systems (NeurIPS 2023).

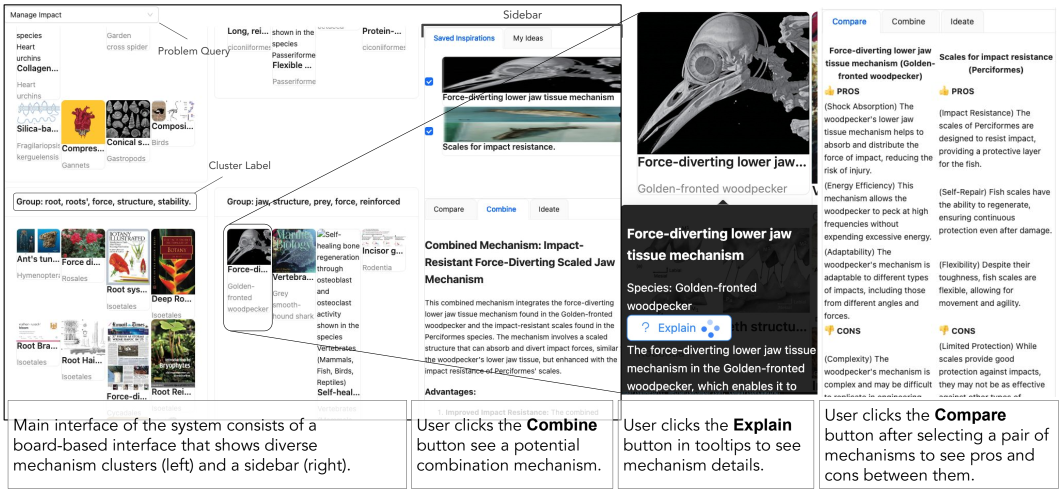
Manage Impact
Problem Query
Sidebar
Cluster Label
Group: root, roots', force, structure, stability.
Group: jaw, structure, prey, force, reinforced
Saved Inspirations
My Ideas
Force-diverting lower jaw tissue mechanism
Scales for impact resistance.
Compare
Combine
Ideate
Force-diverting lower jaw...
Golden-fronted woodpecker
Explain

Combined Mechanism: Impact-Resistant Force-Diverting Scaled Jaw Mechanism

This combined mechanism integrates the force-diverting lower jaw tissue mechanism found in the Golden-fronted woodpecker and the impact-resistant scales found in the Perciformes species. The mechanism involves a scaled structure that can absorb and divert impact forces, similar the woodpecker's lower jaw tissue, but enhanced with the impact resistance of Perciformes' scales.

Advantages:

1. Improved Impact Resistance: The combined

Force-diverting lower jaw tissue mechanism

Species: Golden-fronted woodpecker

The force-diverting lower jaw tissue mechanism in the Golden-fronted woodpecker, which enables it to

Force-diverting lower jaw tissue mechanism (Golden-fronted woodpecker)	Scales for impact resistance (Perciformes)
👍 PROS	👍 PROS
(Shock Absorption) The woodpecker's lower jaw tissue mechanism helps to absorb and distribute the force of impact, reducing the risk of injury.	(Impact Resistance) The scales of Perciformes are designed to resist impact, providing a protective layer for the fish.
(Energy Efficiency) This mechanism allows the woodpecker to peck at high frequencies without expending excessive energy.	(Self-Repair) Fish scales have the ability to regenerate, ensuring continuous protection even after damage.
(Adaptability) The woodpecker's mechanism is adaptable to different types of impacts, including those from different angles and forces.	(Flexibility) Despite their toughness, fish scales are flexible, allowing for movement and agility.
👎 CONS	👎 CONS
(Complexity) The woodpecker's mechanism is complex and may be difficult	(Limited Protection) While scales provide good protection against impacts, they may not be as effective

Main interface of the system consists of a board-based interface that shows diverse mechanism clusters (left) and a sidebar (right).

User clicks the **Combine** button see a potential combination mechanism.

User clicks the **Explain** button in tooltips to see mechanism details.

User clicks the **Compare** button after selecting a pair of mechanisms to see pros and cons between them.

Figure 2: BIOSPARK interface and a subset of available interaction affordances.

limitations can impede the mapping process between biological and engineering systems [11, 21, 9], while the need for situating found analogies within the constraints and scope of the original design problem adds additional complexity for designers developing their own ideas. To mitigate these challenges, we construct a scalable analogy pipeline that balances the cost of expert-curation and the limitations of naïve LLM-based generation, and develop an interactive system that aids designers in recognizing and understanding the relevance and adaptability of analogs for a chosen design problem.

2 System Description, Case Studies, and Future Work

Biological-Analogical Mechanisms Dataset Generation. To source a set of diverse, high-quality biological mechanisms for a given problem, BIOSPARK starts from a seed set of expert-curated biological mechanisms on AskNature (Fig. 1, Step 1). We prompt GPT4 [19] to structure raw AskNature blog posts' text into problem-mechanism schemas (See §B.1 for implementation details). Using each schema as a seed, we iteratively prompt GPT4 to find relevant mechanisms for the given mechanism and problem, using an even mixture of breadth- and depth-focused expansion strategies (Fig. 1, Step 2; §B.2). At each iteration, we construct a taxonomy tree based on organisms represented in the generated mechanisms up to that point, and extract *sibling* taxonomic ranks (for breadth-focused expansion) or *children* taxonomic ranks (for depth-focused expansion), given a reference rank. Extracted taxonomic ranks in the respective expansion strategies are appended to the prompt as constraints (*e.g.,* "come up with a few biological `classes` not in `{...names of excluded classes...}`"). Finally, accompanying representative mechanism images are searched (§B.4).

Interacting with Mechanism Inspirations. The mechanisms dataset is clustered into semantic groups with labels showing the most frequent unigrams appearing in each cluster's mechanism texts, excluding stopwords (Fig. 1, Step 3; §B.3). Users can mouse-over each mechanism to see a detailed explanation or select two mechanisms to compare and combine (see Fig. 2). The interface includes an 'Ideate' tab which brings out a rich text editor in the sidebar where users describe their own ideas based on interacting with mechanism inspirations, and a 'Critique' button to receive feedback on the idea's feasibility and areas to revise (§B.5).

Case Studies. To showcase usage scenarios and demonstrate value of BIOSPARK, we report on two case studies in Appendix A.

Future Work. We end with avenues of future research (Appendix C) for exploration in this area.

3 Conclusion

In this work we present BIOSPARK, an end-to-end system for generating a biological-analogical mechanisms dataset and an interactive interface that facilitates learning new biological mechanisms for a design challenge, and synthesizing new solution ideas inspired by analogical mechanisms.

A Case Studies

A.1 ‘Design a secure bike rack for sedans’

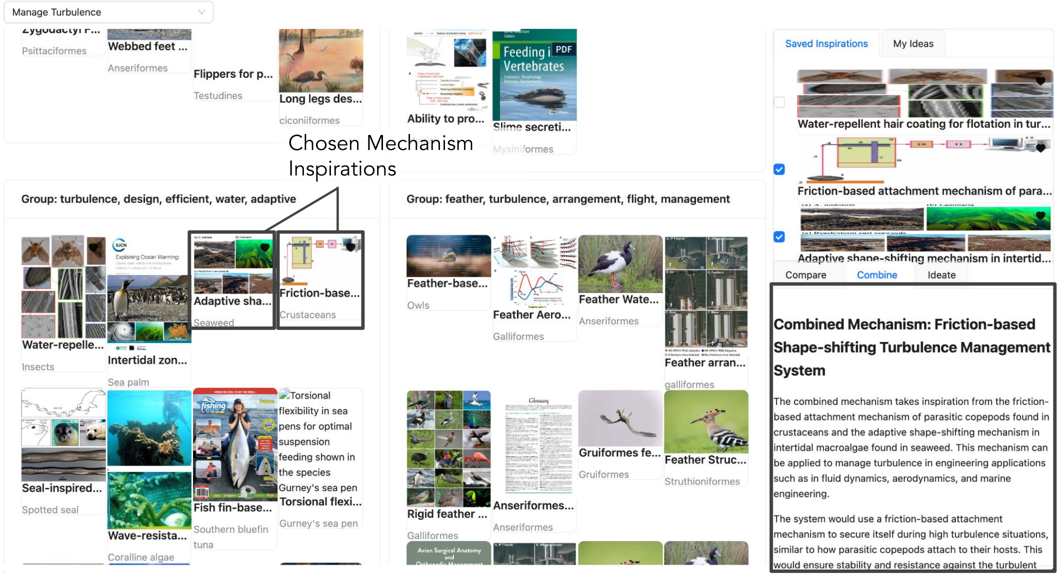

Combined Mechanism: Friction-based Shape-shifting Turbulence Management System

The combined mechanism takes inspiration from the friction-based attachment mechanism of parasitic copepods found in crustaceans and the adaptive shape-shifting mechanism in intertidal macroalgae found in seaweed. This mechanism can be applied to manage turbulence in engineering applications such as in fluid dynamics, aerodynamics, and marine engineering.

The system would use a friction-based attachment mechanism to secure itself during high turbulence situations, similar to how parasitic copepods attach to their hosts. This would ensure stability and resistance against the turbulent

Figure 3: BIOSPARK interface featuring biological mechanisms for the query design challenge ‘Manage Turbulence’. User-chosen mechanisms are: “Adapative shape-shifting mechanism in intertidal microalgae” and “Friction-based attachment mechanism of parasitic copecods.” A new combination idea is shown in the ‘Combine’ tab, facilitating further synthesis and ideation.

Consider an automobile designer who is working on a new bike rack design for sedans. The designer identifies the core challenge of designing a secure bike rack as turbulence management, caused by variable wind resistance due to changing vehicle speeds and additional factors from inclement weather conditions, such as precipitation. When the designer clicks on the ‘Manage Turbulence’ problem as a query (Fig. 3, top right), BIOSPARK displays mechanism clusters related to the problem. She quickly identifies particularly interesting and idea-provoking mechanisms such as “Adapative shape-shifting mechanism in intertidal microalgae” and “Friction-based attachment mechanism of parasitic copecods,” two quite different mechanisms, yet commonly observed in organisms that live underwater. When exposed to turbulence cues, shape-shifting microalgae adaptively reconfigure their surface shapes to reduce resistance from strong currents. This provides her a design idea for adaptively “lying down” bike racks on the roof of the car that reduce air drag, and a “tail basement” bike rack attachment to the back of the car that joins the rotational axis of the front wheel of the bikes and adjusts during turbulence by flexibly (counter-)swinging bikes, mimicking bird tail motions. She clicks on the ‘Combine’ tab with the two mechanisms selected to experiment with design ideas. The result describes an approach that simply combines the friction-based mechanism to the shape-shifting mechanism. This gives her a new idea for further securing the racks, by incorporating designs inspired from the adhesion mechanism of parasitic copecods, such as microridges and channels for increasing friction upon contact.

Interaction demonstration (video): `https://rb.gy/3l97g`

A.2 ‘Design robust collision resistance for racing cars’

For another automobile designer, who is working on improving collision resistance for racing cars, the core challenge might be the management of compression during a collision. When the designer clicks on the ‘Manage Collision’ problem as a query (Fig. 4, top right), BIOSPARK once again displays biological mechanism clusters related to the problem. He finds “Interlocking barbules for structural integrity during flight” and “Pressure management in segmented skeletons,” interesting mechanisms employing different approaches to compression management (the former commonly found in feathers, the latter among isopods). He selects the ‘Combine’ tab to obtain design ideas that could benefit from

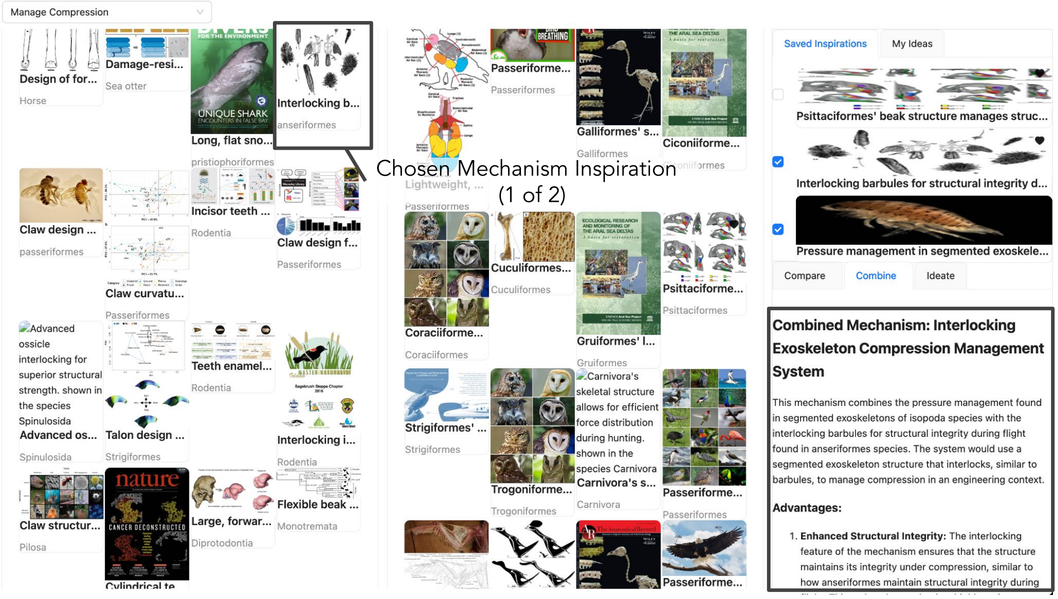

Combined Mechanism: Interlocking Exoskeleton Compression Management System

This mechanism combines the pressure management found in segmented exoskeletons of isopoda species with the interlocking barbules for structural integrity during flight found in anseriformes species. The system would use a segmented exoskeleton structure that interlocks, similar to barbules, to manage compression in an engineering context.

Advantages:

1. **Enhanced Structural Integrity:** The interlocking feature of the mechanism ensures that the structure maintains its integrity under compression, similar to how anseriformes maintain structural integrity during

Figure 4: BioSpark displays biological mechanisms for the query design challenge ‘Manage Compression’. User-chosen mechanisms are: “Interlocking barbules for structural integrity during flight” and “Pressure management in segmented skeletons.”

both mechanisms. The result describes an approach that applies interlocking structures to exoskeleton segments, providing him with insight into a ‘sliding-lock’ mechanism for modular panels of the racing car body. This mechanism allows for extensive compression upon collision, thus increasing impact time (subsequently decreasing peak impact force) and more effectively dissipating force to non-critical body components.

Interaction demonstration (video): `https://rb.gy/2acmp`

B Implementation Details

B.1 Structuring AskNature blog posts’ text into a seed set of problem-mechanism schemas

AskNature.org provides a curated list of organisms with detailed descriptions of their unique strategies to functional problems (*e.g.,* ‘Manage Impact’, ‘Modify Speed’). The organisms and strategies can be grouped by function and viewed as a list. To curate a seed set of high-quality mechanisms, we first choose a functional problem p predicted to be highly relevant to automobile designers, excluding irrelevant functions such as ‘Adapt Behaviors’, ‘Adapt Genotype’, ‘Coevolve’, ‘Maintain Community’ We access the sub-list of organisms $o \in O$ and strategies posted to p on AskNature’s group-by-function page by parsing the `HTML` code using the `BeautifulSoup` package on `Python`. We then access the blog post for each organism-strategy page using the parsed URL and parse the returned `HTML` page to get the title, description, and references (if available).

At this stage, the returned unstructured text is yet to contain a succinct mechanism description. Furthermore, we found that some blog posts do not contain any body text despite having a title and are accessible via the URL. Some of these missing blog posts indicated that they are in maintenance and/or planned to be updated. To structure the raw blog post text AskNature$_{(o,p)}$, we prompt GPT4 to succinctly describe (*i.e.,* using 12 words or less) the core mechanism (*i.e.,* excluding the qualities or effects, and focusing on mechanisms with engineering design implications), given (o, p) (if blog post text is missing) or $(o, p, \text{AskNature}_{(o,p)})$. The returned mechanism description m along with the function description makes up the problem-mechanism schema for each organism: $\{o \in O | (p, m, o)\}$.

B.2 Iteratively constructing taxonomic trees with organisms represented in found mechanisms and performing structured expansion of a tree

To enable structured diversification of organisms and their mechanisms beyond prior work that relied on token-level manipulation or naïvely prompting LLMs, we guide LLMs *how* and *where* to expand by leveraging organism taxonomic hierarchies. At each iteration of expansion (Fig. 1, Step 2), we aggregate the organisms represented in found mechanisms up to that point, and construct a taxonomic tree featuring seven levels of hierarchy on Tree of Life: {`domain, kingdom, phylum, class, order, family, genus, species`}, where `domain` representing the highest level and `species` representing the lowest level on the hierarchy.

Given this tree, we aim to identify sparsely populated branches for expansion. We cut the tree at a given reference expansion level (*e.g.,* `class`), and sort the taxonomic ranks (nodes) on that level by the number of its immediate children nodes[1], in an increasing order. For performance, we batch 10 prompts to send to GPT4 for expansion. For half of the prompts, we instruct **breadth-first expansion** which asks GPT4 to first identify sibling nodes at the given reference taxon level and existing nodes (up to 50 most populated nodes). For example, the prompt asks "come up with a few biological `classes` not in {`...names of excluded classes...`}". The breadth-first expansion prompt then instructs GPT4 to repeat the following: 1) Choose one taxon from the list it came up with; 2) Succinctly describe (*i.e.,* using 14 words or less) new mechanisms m related to a problem p. For the rest of the prompts, we instruct **depth-first expansion** which asks GPT4 to first identify a new children node at the given reference taxon level and existing children nodes (up to 50 randomly sampled children). For example, the prompt asks "come up with a few biological `families` in `order araneae` that are not any of {`araneidae`, ...}". The depth-first expansion prompt then instructs GPT4 to repeat a similar procedure as breath-first expansion. In the prototype system, we run 10 batches for expansion to construct dataset of mechanisms for each problem.

The returned list of mechanisms and organisms text are then fed into the second GPT4 prompt for structuring them into a list of {`mechanism, organism`} dictionaries. Finally, using each organism name, we prompt GPT3.5-turbo to retrieve the seven-level taxonomic hierarchy. We use GPT3.5-turbo based on our evaluation of its comparable accuracy levels to GPT4 for this task (the highest rate of misaglignment occured in the `family` rank, with a rate of 6.25% negative cases, which is equivalent to 15 out of 240 organisms, that differed from the gold-standard answers extracted from corresponding Wikipedia pages' scientific classification table).

B.3 Semantic Clustering of Mechanisms

The mechanism descriptions are clustered in semantically related groups for designers' efficient viewing. We first embed the mechanism descriptions for a given problem into high-dimensional vectors using the `text-embedding-ada-002` end-point of the OpenAI API. We group the embeddings into $k = 20$ clusters using k-Means clustering. Each group label is generated by chunking the mechanism descriptions in each cluster into lower-cased words, excluding stopwords. We count the frequency of each word and take the top 5 most frequent words as a description of each cluster.

B.4 Representative Mechanism Image Curation

To aid designers' visual understanding of and pique curiosity for biological-analogical mechanisms, we retrieve representative images for corresponding textual mechanism descriptions. We use Google Custom Search[2] with queries as "`[organism name]:[mechanism description]`" and the file type set to images and the safe search mode enabled. We choose the first place result of Custom Search as the visual representation of each mechanism.

B.5 Interaction Features: Explain, Compare, Combine, and Critique

To facilitate designers' understanding and synthesis of mechanism inspirations, we develop several interaction features available on the interface (Fig. 2). The **Explain** button is located in tooltips that pop up when the user places the mouse over on a mechanism card in the board UI (Fig. 2, first panel).

[1] alternatively, the entire size of the subtree, rather than immediate children, could be used for sorting

[2] `https://developers.google.com/custom-search/v1/overview`

When the user clicks on the button, BIOSPARK sends a prompt to GPT4 requesting elaboration of the interacted mechanism and the organism in the context of the chosen engineering design problem. The **Compare** tab is located in the control bar of the sidebar of the interface. To use this, users need to first click on (at least) two mechanism cards from the left, saving them to the 'saved inspirations' panel at the top of the sidebar. There, users can check any two of the saved mechanisms they wish to compare. BIOSPARK sends a prompt to GPT4 when the user clicks on the tab, requesting comparison of pros and cons between the two mechanisms in the context of the chosen engineering problem. The result is formatted into a markdown table, with each mechanism as the header followed by pros and cons rows detailing each point. The **Combine** tab is also located in the control bar of the sidebar in the interface. Similarly with Compare, users can check two saved mechanisms they wish to see combined. BIOSPARK sends a prompt to GPT4 then requesting elaboration of a mechanism that combines the two selected mechanisms, and explain its potential advantages in the context of the chosen engineering problem. The result is also formatted into a markdown page using section title and headers for demarcating the content. Finally, the **Critique** button is located inside the Ideate tab. Upon clicking the Ideate tab, users can type in their own idea in the rich text editor in the opened tab, and optionally click on the button below to receive constructive feedback on it. BIOSPARK sends a prompt to GPT4 with the content of the text editor describing the idea, and requests additional revision that may improve the quality, such as anticipated failure modes and potential improvements.

C Discussion and Future Work

C.1 Closing the 'Gulf of Geometric Transfer' with combinatoric visual mechanism generation

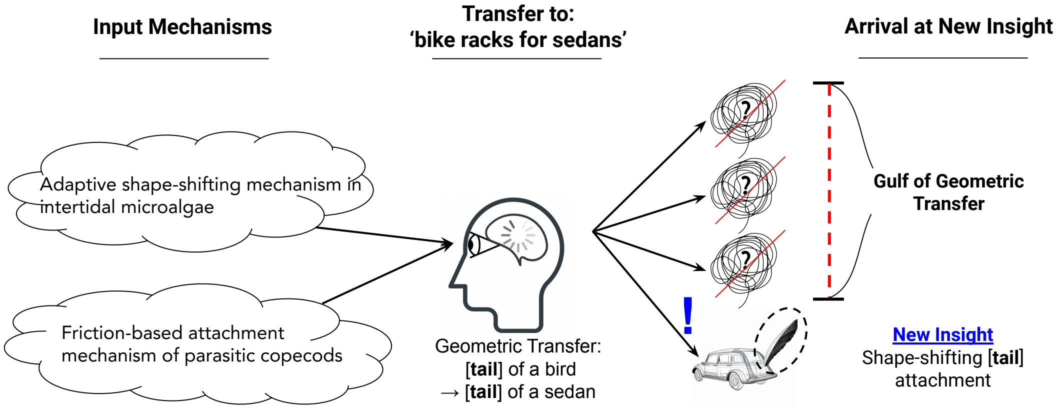
Input Mechanisms
Transfer to:
'bike racks for sedans'
Arrival at New Insight
Adaptive shape-shifting mechanism in intertidal microalgae
Friction-based attachment mechanism of parasitic copecods
Geometric Transfer:
[tail] of a bird
→ [tail] of a sedan
Gulf of Geometric Transfer
!
New Insight
Shape-shifting [tail] attachment

Figure 5: When the designer arrives at the new insight of designing a tail-like, shape-shifting bike rack attachment (Appendix A.1), she is combining two source mechanisms and transferring the new idea into a new geometry in the target domain ('at the *tail* of a sedan') that could suit the combined mechanism. The successful generation of this combination comes after a series of failed attempts to generate geometries, which we call the **Gulf of Geometric Transfer**.

The designer in Case Study 1 (Appendix A.1) arrives at a successful transfer of mechanisms naturally after several unsuccessful iterations of generation. A key process in transferring a new idea into the target domain involves visualizing a candidate geometry, such as in this case 'a bike rack attachment to the *tail* of a sedan'. The process of identifying an effective geometry for transfer like 'at the tail of a sedan' can resemble a combinatoric search, where the designer visualizes the location of integration for the new idea in the target object, from a possible source object and its geometry, such as 'the tail of a bird'. The process of evaluating an effective geometry for transfer may further involve analogical processing, for instance, by comparing the integrated geometry to familiar objects and judging the aesthetic and functional value of such a geometry.

We hypothesize that one challenge of this transfer process is what we term '**Gulf of Geometric Transfer**' (see Fig. 5, right; the concepts related to this namesake in design and usability studies are the Gulf of Execution and the Gulf of Evaluation [18]), which refers to the failed iterations that one has to endure before obtaining a successful geometric insight that supports the new idea. To bridge this gulf, we envision a future work area that aids designers in visualizing what successful geometries that facilitate transfer of ideas might look like, and also helps them eliminate any blind spots in the process that could hinder new insights. A fruitful starting point to this end may involve generating visual representations of geometries that are related to the new abstract idea from various key orientations, such as {`up`, `down`, `front`, `rear`, `left`, `right`, `internal`, `external`}. Enumerating these geometries visually could assist designers in effectively reviewing possible geometries and facilitating their imagination for alternative geometries that may be useful for further consideration.

C.2 Supporting Generative 'Frame Slips'

The geometric transfer discussed in the previous section can be understood as one form of a more generalized phenomenon in analogical transfer, where a conceptual frame 'slips' into a different one, thereby generating novel insight. For example, Niels Bohr transferred the concept of planets orbiting the sun to the domain of atomic structure to propose a new model in which electrons orbit the nucleus in a similar manner [9]. This new structure of the atom significantly differed from the dominant 'plum-pudding' model of the time, leading to a *re-structuring* of the concept [10]. Here, *re-structuring* is related to conceptual frames that successfully 'slip' into a new, analogous problem domain, causing structural changes in the conceptaul framework of the receiving domain or vice versa. In the case studies, the concept of the 'tail of a car' slipped into geometrically and syntactically similar 'tail of a bird', engendering a new insight. Similarly while studying how crowds generate design concepts for children's chairs, Yu et al. [27] found that the concept of a playground 'slide' slipped into another semantically related playground concept 'swing set', which then gave a novel design idea of swing chairs.

Understanding the process of '**Frame Slipping**' is therefore crucial for improving analogical transfer. Existing examples of 'frame slips' imply that such transitions can occur through different relational paths, for example, syntactic, semantic, functional, abstracted commonality, geometric, or other means. While geometric transfer may impose a more concrete form of transformation onto the transfer process, other means of transfer may be less understood, possibly involving more complex multimodal processes that combine multiple relational paths at once. Nonetheless, significant value of an inspiration often results from enabling frame slips to new parts of the design space, opening up not only a single point of an idea, but an entire space of ideas for future exploration.

C.3 Supporting any problem queries by incorporating rich problem-mechanism relations

One of BIOSPARK's core trade-offs is its limited querying flexibility. Though the five pre-generated problem queries provide a useful entry to mechanism organisms that may be applicable to a diverse set of design challenges, it comes at the cost of an inability to query biological mechanisms for *any* engineering design problems. As designers and engineers progress in interacting with the system, they may naturally come up with follow-up queries that may differ from the source queries, that may emphasize important constraints around the design problem, or specify low-level details newly understood to be important to consider. In order to enable search by free-text problem queries, the underlying data model needs to be extended to contain multiple problem-mechanism relations beyond the single problem present in the schema $\{\forall i | (p_i, m_i, o_i)\}$, and into an enriched dataset with mappings between problems $p_1, p_2, \cdots, p_n \in P$ and an applicable mechanism m_k, as commonly the case in engineering (*e.g.,* 'spider silk' can be used for multiple engineering challenges such as replacing steel bars in concrete or wound suture and prosthesis [14]).

One way to expand the rich problem-mechanism relations in a scalable manner is to prompt LLMs to come up possible engineering design problems that a given biological mechanism could be applied to. Here, naïve prompting may suffer from conceptual redundancy, analogous to the challenge of curating diverse mechanisms, that limits the diversity in mechanism-to-problem mappings. Another approach may be to intelligently use the existing dataset $\{\forall i | (p_i, m_i, o_i)\}$ to identify similar mechanisms m_i in (p_i, m_i, o_i) and m_j in (p_j, m_j, o_j) that can be mapped onto disparate problems: $m_i \sim m_j \to p_i, p_j$. This approach however assumes the presennce of many correlated such mechanisms with disparate problem pairs in the dataset, which need empirical examination for support. Once the enriched dataset $\{\forall i | m_i \to S(m_i)\}$ (Here, $S(m_i)$ denotes the set of engineering design problems that m_i is applicable to) is made available, one simple approach for allowing querying on any problem text is to construct a similarity search index (*e.g.,* the HNSW index of FAISS [16]) using a chosen text embedding approach.

C.4 Supporting improved visual representation of biological mechanisms

Future work may focus on promising methodologies for improved visual representation of biological mechanisms. Augmentative approaches to the simple retrieval approach described in the current prototype include 'Retrieve and Filter' and 'Highlight Species Subparts'. The 'Retrieve and Filter' method involves a process of purging non-representative and low-quality images and ranking the remaining ones based on their semantic distance to the mechanism text, resulting in images with higher semantic coherence with the corresponding mechanism description. On the other hand, the 'Highlight Species Subparts' method will involve sourcing for representative species illustrations or photographs, and highlighting components corresponding to mechanisms via a translucent overlay. This method can be applied when previous method fails to yield high-quality images.

While 'Highlight Species Subparts' presents potential, we also anticipate certain limitations with it. For example, specific mechanisms could only apply to a tiny animal area, rendering a whole body image inappropriate, such as in cases of the adhesive paws of geckos or the hierarchical layers in pomelo peels. Additionally, in certain cases, embedding textual labels within a diagrammatic representation, as opposed to using realistic species images, could provide more beneficial insights, annotating key components rather than merely a visual highlight. Understanding the trade-off between these strategies could be beneficial.

Alternatively, generative AI could unlock new approaches to generate representative mechanism images on demand. Here, a systematic method such as focusing on cases where retrieval-based images fail could prove effective. A practical approach would involve creating a mechanism taxonomy potentially in tabular format showcasing a mechanism's text, its type, and specific example images that unsuccessful retrieval-based approaches attempted to represent. This could offer valuable insights into the areas and reasons why our retrieval-based imaging models are not successful, thereby guiding the further development of generatively-focused models.

Acknowledgments and Disclosure of Funding

This research was supported by the Toyota Research Institute.

References

[1] Jinheon Baek, Alham Aji, and Amir Saffari. Knowledge-augmented language model prompting for zero-shot knowledge graph question answering. In *Proceedings of the First Workshop on Matching From Unstructured and Structured Data (MATCHING 2023)*, pages 70–98, Toronto, ON, Canada, July 2023. Association for Computational Linguistics.

[2] Janine M Benyus. *Biomimicry: Innovation inspired by nature*. Morrow New York, 1997.

[3] Chris Broeckhoven and Anton du Plessis. Escaping the labyrinth of bioinspiration: biodiversity as key to successful product innovation. *Advanced Functional Materials*, 32(18):2110235, 2022.

[4] Amaresh Chakrabarti, Prabir Sarkar, B Leelavathamma, and BS Nataraju. A functional representation for aiding biomimetic and artificial inspiration of new ideas. *Ai Edam*, 19(2):113–132, 2005.

[5] Hyunmin Cheong and LH Shu. Retrieving causally related functions from natural-language text for biomimetic design. *Journal of Mechanical Design*, 136(8):081008, 2014.

[6] John Chung, Ece Kamar, and Saleema Amershi. Increasing diversity while maintaining accuracy: Text data generation with large language models and human interventions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 575–593, Toronto, Canada, July 2023. Association for Computational Linguistics.

[7] Jon-Michael Deldin and Megan Schuknecht. The asknature database: enabling solutions in biomimetic design. In *Biologically inspired design: computational methods and tools*, pages 17–27. Springer, 2013.

[8] D. Gentner, S. Brem, R. W. Ferguson, P. Wolff, A. B. Markman, and K. D. Forbus. Analogy and Creativity in the Works of Johannes Kepler. In T. B. Ward, J. Vaid, and S. M. Smith, editors, *Creative thought: An investigation of conceptual structures and processes*, pages 403–459. American Psychological Association, Washington D.C., 1997.

[9] Dedre Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive science*, 7(2):155–170, 1983.

[10] Dedre Gentner. Analogy in scientific discovery: The case of johannes kepler. In *Model-based reasoning: Science, technology, values*, pages 21–39. Springer, 2002.

[11] Dedre Gentner and Russell Landers. Analogical reminding: A good match is hard to find. In *Unknown Host Publication Title*, pages 607–613. IEEE, 1985.

[12] Mary L Gick and Keith J Holyoak. Analogical problem solving. *Cognitive psychology*, 12(3):306–355, 1980.

[13] Ashok K Goel, Swaroop Vattam, Bryan Wiltgen, and Michael Helms. Cognitive, collaborative, conceptual and creative—four characteristics of the next generation of knowledge-based cad systems: A study in biologically inspired design. *Computer-Aided Design*, 44(10):879–900, 2012.

[14] Yunqing Gu, Lingzhi Yu, Jiegang Mou, Denghao Wu, Peijian Zhou, and Maosen Xu. Mechanical properties and application analysis of spider silk bionic material. *e-Polymers*, 20(1):443–457, 2020.

[15] Tom Hope, Joel Chan, Aniket Kittur, and Dafna Shahaf. Accelerating innovation through analogy mining. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, pages 235–243, New York, NY, USA, 2017. ACM.

[16] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547, 2019.

[17] Hyeonsu B. Kang, Xin Qian, Tom Hope, Dafna Shahaf, Joel Chan, and Aniket Kittur. Augmenting scientific creativity with an analogical search engine. *ACM Trans. Comput.-Hum. Interact.*, 29(6), nov 2022.

[18] Donald A. Norman and Stephen W. Draper. *User Centered System Design; New Perspectives on Human-Computer Interaction*. L. Erlbaum Associates Inc., USA, 1986.

[19] OpenAI. Gpt-4 technical report, 2023.

[20] R. Oppenheimer. Analogy in science. *American Psychologist*, 11(3):127–135, 1956.

[21] Brian H Ross. This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13(4):629, 1987.

[22] LH Shu, K Ueda, I Chiu, and H Cheong. Biologically inspired design. *CIRP annals*, 60(2):673–693, 2011.

[23] Dennis Vandevenne, Thomas Pieters, and Joost R Duflou. Enhancing novelty with knowledge-based support for biologically-inspired design. *Design Studies*, 46:152–173, 2016.

[24] Swaroop Vattam and Ashok Goel. 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*, 2013.

[25] Swaroop S Vattam and Ashok K Goel. 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, pages 177–186, 2011.

[26] Julian FV Vincent and Darrell L Mann. Systematic technology transfer from biology to engineering. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 360(1791):159–173, 2002.

[27] Lixiu Yu, Robert E. Kraut, and Aniket Kittur. Distributed analogical idea generation with multiple constraints. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, CSCW '16, page 1236–1245, New York, NY, USA, 2016. Association for Computing Machinery.