

LLM2TEA: An Agentic AI Designer for Discovery with Generative Evolutionary Multitasking

Melvin Wong¹, Jiao Liu¹, Thiago Rios², Stefan Menzel², Yew Soon Ong^{1,3}

¹College of Computing & Data Science (CCDS), Nanyang Technological University (NTU), Singapore

²Honda Research Institute Europe (HRI-EU), Offenbach am Main, Germany

³Centre for Frontier AI Research (CFAR), Agency for Science, Technology and Research (A*STAR), Singapore

{wong1357, jiao.liu, asysong}@ntu.edu.sg, {thiago.rios, stefan.menzel}@honda-ri.de

Abstract—This paper presents LLM2TEA, a Large Language Model (LLM) driven MultiTask Evolutionary Algorithm, representing the first agentic AI designer of its kind operating with generative evolutionary multitasking (GEM). LLM2TEA enables the crossbreeding of solutions from multiple domains, fostering novel solutions that transcend disciplinary boundaries. Of particular interest is the ability to discover designs that are both novel and conforming to real-world physical specifications. LLM2TEA comprises an LLM to generate genotype samples from text prompts describing target objects, a text-to-3D generative model to produce corresponding phenotypes, a classifier to interpret its semantic representations, and a computational simulator to assess its physical properties. Novel LLM-based multitask evolutionary operators are introduced to guide the search towards high-performing, practically viable designs. Experimental results in conceptual design optimization validate the effectiveness of LLM2TEA, showing 97% to 174% improvements in the diversity of novel designs over the current text-to-3D baseline. Moreover, over 73% of the generated designs outperform the top 1% of designs produced by the text-to-3D baseline in terms of physical performance. The designs produced by LLM2TEA are not only aesthetically creative but also functional in real-world contexts. Several of these designs have been successfully 3D printed, demonstrating the ability of our approach to transform AI-generated outputs into tangible, physical designs. These designs underscore the potential of LLM2TEA as a powerful tool for complex design optimization and discovery, capable of producing novel and physically viable designs.

Index Terms—Agentic AI Designer, Large Language Models, text-to-3D Generative AI, Generative Evolutionary Multitasking, Engineering Design Optimization and Discovery, LLM Evolutionary Optimizers, Evolutionary Multitasking

I. INTRODUCTION

Generative Artificial Intelligence (GenAI) has demonstrated remarkable capabilities in producing diverse creative designs from free-form natural language prompts [1]. Notable examples include the synthesis of creative texts, images, audio, and videos through tools such as SARD [2], Stable Diffusion [3], Make-An-Audio [4], and Veo 2.0 [5]. Over the past few years, a surge of research in GenAI has focused

on enabling the production of high-quality digital media that closely aligns with user preferences through the use of tailored prompts. Despite these widely publicized advancements in digital media, a notable gap remains in applying GenAI to complex engineering design, with comparatively fewer reports addressing its potential within the science and engineering domains.

In the realm of traditional engineering design, which is the primary focus of this work, the promising potential of GenAI in conceptual aerodynamic design has been demonstrated [6], [7]. The major challenge identified in applying GenAI to science and engineering lies not only in generating digital designs exhibiting high novelties, but also in ensuring that the designs can translate into tangible physical solutions that adhere to natural principles [8] or satisfy specific problem contexts and design specifications [9]. Our early research studies have attempted to address this challenge by advancing GenAI for engineering design optimization [10]–[12]. Preliminary results leveraging conventional optimization strategies indicate encouraging progress, particularly through the use of text prompt search to improve aerodynamic performance [6] while maintaining semantic conformity [7]. These efforts underscore the potential of GenAI to bridge the gap between digital innovation and practical applicability, particularly in engineering disciplines. However, navigating the very high-dimensional language search space often demands an integration of symbolic, statistical, and experiential knowledge to support broader generalization and knowledge transfer. This imperative highlights the need for agentic AI approaches that can engage in objective-directed reasoning and guidance as well as adaptive problem-solving.

To this end, a large language model (LLM) driven MultiTask Evolutionary Algorithm (LLM2TEA), is introduced as the agentic AI designer implemented with generative evolutionary multitasking (GEM), structured to facilitate the cross-pollination of solutions across multiple domains. The objective of LLM2TEA is the discovery of novel solutions that transcend disciplinary boundaries while adhering to natural principles and meeting user specifications. LLM2TEA is composed of an LLM-guided multitask evolutionary algorithm, a text-to-3D generative model responsible for phenotype production, a classifier for interpreting semantic representations, and a

physics simulation model to evaluate the physical properties of the generated designs. By solving multiple tasks simultaneously, where each task corresponds to a single-domain search, the exploration of design space is expanded beyond isolated domains and into their intersections, allowing for the emergence of cross-domain designs enabled by knowledge transfer.

Several novel LLM-based multitask evolutionary operators are introduced, incorporating new instruction sets to guide the evolutionary processes. This includes an initialization operator for generating an initial population of genotypes and multitask mating operators, such as crossover, mutation, and selection, to enable the efficient exchange of traits across domains as evolution progresses with the generation of offspring. The resulting genotype population serves as input to the text-to-3D generative model for synthesizing phenotypes. The survival of the best-performing designs is evaluated based on the practical viability of these phenotypes, considering criteria that include physical properties and semantic meanings. In summary, the contributions of this paper are as follows:

- This paper introduces LLM2TEA, a novel agentic AI designer with a generative evolutionary multitasking (GEM) that promotes the crossbreeding of designs from multiple domains, leading to novel solutions that transcend individual domains.
- LLM2TEA is an LLM-driven multitask evolutionary algorithm with operators synthesizing creative text prompts as guides to an auxiliary text-to-3D generative model in producing corresponding novel 3D designs. LLM2TEA represents the first method operating with the GEM paradigm to leverage the capabilities of LLMs for complex design optimization and the discovery of novel digital designs. These digital designs not only exhibit high novelty but are also translatable into tangible physical solutions that conform to natural laws and fulfill specific problem contexts and design specifications.
- A series of complex design and discovery tasks, modeled after real-world design optimization scenarios, is conducted to assess the efficacy of the proposed LLM2TEA. Experimental results indicate that LLM2TEA consistently outperforms the baseline text-to-3D generative model in uncovering novel designs and/or enhancing aerodynamic performance. The proposed novelty score metric is employed for evaluation, with findings showing that over 45% of generated designs are identified as novel, exceeding the mean performance of the baseline.
- It is demonstrated that several designs discovered through LLM2TEA are physically realizable and successfully 3D printed using off-the-shelf stereolithography (SLA) technology. As previously noted, direct outputs from baseline text-to-3D generative models often result in designs that are highly fragmented for physical fabrication. In contrast, the findings indicate that LLM2TEA is capable of identifying physically feasible designs suitable for SLA printing, highlighting its potential as a viable approach for practical rapid design prototyping.

The remainder of the paper is organized as follows: Section

II reviews related work from the literature and provides an overview of key background concepts. Section III introduces LLM2TEA in detail. Section IV presents the experimental setup, and the results and discussions are detailed in Section V. Finally, Section VI concludes this paper, highlighting some potential future directions.

II. RELATED WORKS

This section presents an overview of related work. First, the applications of *text-to-X* generative models (where *X* represents various outputs such as languages, images, or 3D shapes) are reviewed in the context of creative and engineering conceptual design. Next, prompt engineering and search techniques are examined, with a focus on methods to generate effective prompts for text-to-*X* generative models. Finally, recent advancements in the evolutionary multitasking (EMT) paradigm are discussed, with an emphasis on its potential to address cross-domain optimization tasks simultaneously.

A. Text-to-*X* Generative Models in Creative & Engineering Conceptual Design

Recently, text-to-*X* generative models have been introduced to allow users to specify design specifications using free-form natural language prompts. Pioneering work investigated the use of text-to-image generative models to synthesize novel creative digital artifacts that satisfy user preferences [1]. On the other hand, Point-E [13] and Shap-E [14] were among the first models to generate 3D artifacts from text prompts. In contrast to Point-E, which is based on 3D point cloud representations, Shap-E includes additional features, such as a differentiable implementation of the marching cubes 33 algorithm [15] to synthesize 3D polygonal meshes, which are commonly used representation format in computer graphics or for computational simulation-based applications.

B. Prompt Engineering and Search

It has been observed that prompt augmentation significantly influences the feasibility and quality of artifacts generated by text-to-*X* generative models. Consequently, designing effective prompts is crucial for fully harnessing the full potential of these generative models for solving diverse tasks. Several studies have proposed prompting techniques based on handcrafted prompt templates to better align the outputs of the model with target objectives [16], [17]. In contrast, prompt search methods have considered optimization approaches to finding optimal prompts through iterative refinement [18]–[20].

In recent years, prompt search techniques have also been applied in engineering design optimization [6]. Rios et al. proposed a fully automated evolutionary design optimization framework incorporating evolutionary strategies with performance-based guidance to seek for novel 3D car designs using the Shap-E text-to-3D generative model. However, searching within the text-prompt space can often lead to ill-defined design outputs, as Shap-E may not be explicitly trained on a dataset focused on car shapes. To address this limitation, vision-language models [21], [22] were introduced into the

framework to penalize semantically incoherent or structurally implausible designs [7]. The research demonstrated the feasibility of searching effective text prompts to generate designs that meet both practical constraints and visual requirements with evolutionary approaches. However, the resulting design often retained strong resemblances to conventional car shapes, since the optimization process was narrowly targeted at car design objectives.

Due to recent advancements in language comprehension and generation capabilities, LLMs have become effective tools not only for generating prompts but also for serving as optimizers. For instance, Yang *et al.* [23] directly utilized an LLM as the optimizer in the prompt optimization task, demonstrating both the feasibility and effectiveness of such an approach. Guo *et al.* [24] combined evolutionary algorithms with an LLM to enhance prompt optimization, leveraging the strengths of both methods. However, most existing approaches primarily focus on optimizing prompts for a single task. As discussed in Section I, it is believed that incorporating cross-domain learning into the design optimization process can identify prompts that lead to novel design solutions. Therefore, prompt search via evolutionary multitasking techniques is proposed. Such an approach leverages knowledge from multiple domains to explore a broader design space that extends beyond individual domains and embraces its overlap, prompting the discovery of novel cross-domain designs through knowledge transfer.

C. Evolutionary Multitasking

Recently, the evolutionary multitasking (EMT) paradigm has emerged as a population-based search methodology designed to handle multiple optimization tasks concurrently [25]. In contrast to traditional evolutionary algorithms, EMT exploits latent synergies between distinct yet related optimization tasks, thereby enhancing convergence performance. Over the past few years, several successful applications of EMT have been documented across various domains, including manufacturing process design [26], mobile edge computing [27], and vehicle routing [28], [29]. These examples highlight the potential of EMTs for efficiently solving complex, multifaceted problems by leveraging cross-task knowledge sharing.

While there is surging research in using LLMs as evolutionary optimizers to address both traditional optimization [30], [31] and prompt optimization tasks [24], no documented studies have explored exploiting the contrasting characteristics of multiple domains to enable crossbreeding of solutions with evolutionary multitasking, resulting in the discovery of novel artifacts with hybrid features. This paper aims to investigate this possibility, proposing a novel prompt search technique that operates with generative evolutionary multitasking (GEM). The proposed method not only demonstrates effective optimization but also promotes the discovery of novel solutions.

III. AGENTIC AI DESIGNER WITH GENERATIVE EVOLUTIONARY MULTITASKING (GEM)

In this section, the proposed LLM2TEA is presented. First, the multitask optimization problem of interest is described. Next, each distinctive feature of the LLM2TEA, describing

its role in the multitask evolutionary algorithm, is examined in detail.

A. Preliminaries

Let $\mathcal{Z} \subseteq \mathbb{Z}_{\geq 0}^{L \times V}$ denote the search space of sequential tokens, where L is the maximum sequence tokens length and V is the vocabulary size. Given a set of tasks $\{T_1, \dots, T_K\}$, the LLM constructs a genotype (prompt) ($\mathbf{z}_k \in \mathcal{Z}$) relevant to the k -th task using the following template:

$$\mathbf{z}_k = \text{tokenize}(\text{"A } \langle \mathbf{t}_k \rangle \text{ in the shape of } \langle \boldsymbol{\rho}_k \rangle \text{"})$$

Here, a user provides a *target domain* " \mathbf{t}_k " that defines the task of interest (e.g., " \mathbf{t}_k " as "car", "airplane" or others). The variable $\boldsymbol{\rho}_k$ represents the part of the prompt that LLM2TEA will optimize. Based on this setting, the single-objective function for the k -th task is formulated as follows:

$$\begin{aligned} & \forall T_k, k \in \{1, \dots, K\}, \\ & \min_{\mathbf{z}_k \in \mathcal{Z}} : \mathbb{E}_{p_\theta(\mathbf{x}|\mathbf{z}_k)} [(1 - \alpha)f_{phy}(\mathbf{x}) + \alpha(1 - f_{visual}(\mathbf{x}|\mathbf{t}_k))], \end{aligned} \quad (1)$$

where $p_\theta(\mathbf{x}|\mathbf{z}_i)$ represents a *text-to-3D* model that generates the phenotypes (3D shapes) \mathbf{x} based on the prompt \mathbf{z}_k . The function $f_{phy}(\cdot)$ is an evaluator that assesses the shape \mathbf{x} against a target physical criterion. Additionally, LLM2TEA visually assesses the shape \mathbf{x} with a vision-language model $f_{visual}(\cdot|\mathbf{t}_i)$ by measuring the probability that the generated 3D design visually aligns with the target domain \mathbf{t}_k . The α hyperparameter is introduced as a predefined weight to balance the visual fitness score versus the physical fitness score.

Correspondingly, a multi-objective optimization problem with M objectives can be formulated as follows:

$$\begin{aligned} & \forall T_k, k \in \{1, \dots, K\}, \\ & \min_{\mathbf{z}_k \in \mathcal{Z}} : \left\{ \mathbb{E}_{p_\theta(\mathbf{x}|\mathbf{z}_k)} [f_{phy}^{(1)}(\mathbf{x})], \dots, \mathbb{E}_{p_\theta(\mathbf{x}|\mathbf{z}_k)} [f_{phy}^{(M)}(\mathbf{x})] \right\} \quad (2) \\ & \text{s.t. } \underline{c} \leq f_{visual}(\mathbf{x}|\mathbf{t}_k) \leq \bar{c}, \end{aligned}$$

where \underline{c} and \bar{c} control the exploration of designs within and near the domain boundary, and solutions trade-offs among the multiple criteria are of interest. The setting of these objective functions ensures that the optimized prompt generates shapes that not only meet the target physical criteria but also resemble the conceptual elements of the target domain represented by \mathbf{t}_k .

B. LLM2TEA Overview

Given the very high-dimensional nature of the language search space, assume the existence of a function $g : \mathbb{Z}_{(C \times V_c)^H} \rightarrow \mathbb{Z}^{L \times V}$ that can construct the prompt \mathbf{z}_k , where H denotes the historical context, C is the maximum sequence context length and V_c is the context vocabulary size. Further, the hidden Markov property is assumed to hold, which states that all necessary information to derive an outcome exists in the current state (or context in our case), thereby reducing $H = 1$. Under these assumptions, the capabilities of an LLM $p_{llm}(\mathbf{z}_k|\mathbf{t}_k)$ are utilized as the function g to construct

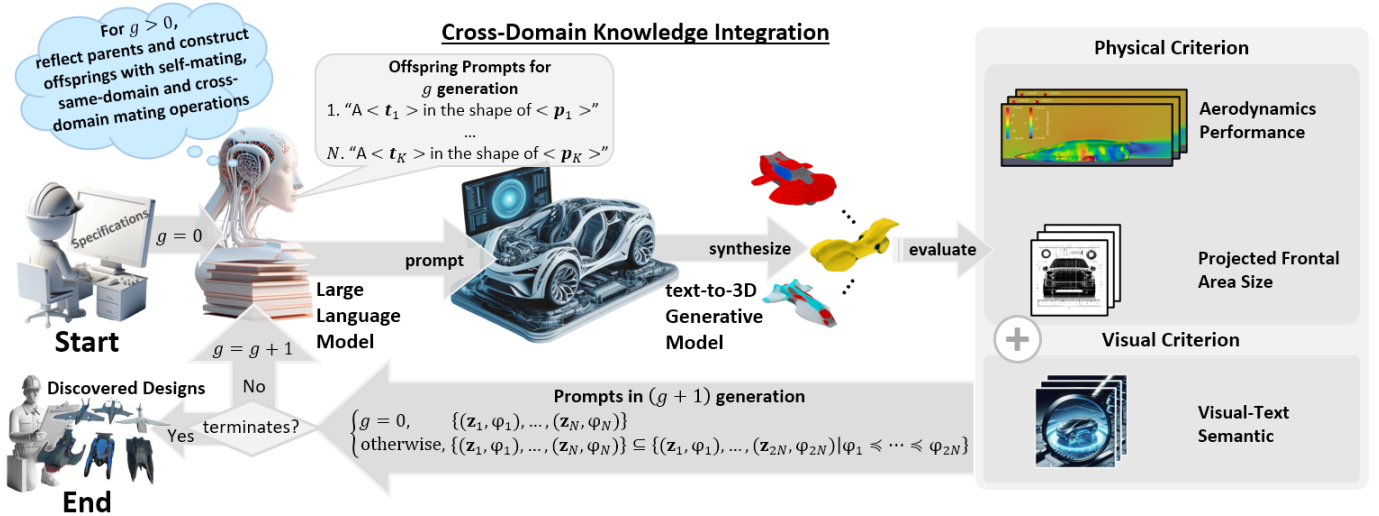


Fig. 1. LLM2TEA is a novel LLM-driven MultiTask Evolutionary Algorithm for complex design and discovery. LLM2TEA optimizes multiple tasks simultaneously, leveraging cross-domain knowledge learning to drive LLM in evolving prompts. The prompts serve as guides for an auxiliary text-to-3D generative model, expanding exploration into cross-domain regions.

effective prompts that guide the text-to-3D model in generating aerodynamically performing vehicle designs. With these considerations, LLM2TEA (Fig. 1) comprises an LLM that constructs prompts from a set of instructions (depicted in Fig. 2) containing the user specifications embedded in z_k . The generated prompts are passed to a text-to-3D generative model that synthesizes designs, which are then evaluated against a set of criteria to compute the fitness with respect to the target objective. Details of the key algorithmic features of LLM2TEA for various proposed approaches in complex design and discovery are elaborated upon in this section.

1) *Population Initialization*: LLM2TEA begins with instructing an LLM to generate an initial population of genotypes, each genotype being a text prompt based on free-form natural language as discussed in Section III-A. Specifically, the LLM is presented with a description of the optimization problem (shown in Fig. 2(a)) and is instructed to generate a population of genotypes (shown in Fig. 2(b)). These genotypes are then passed to a text-to-3D generative model p_θ for synthesizing the corresponding 3D designs.

2) *Offspring Generation*: It has been demonstrated that the in-context learning capability of LLM can incorporate historical objective evaluation data to enhance optimization guidance [30]. Building on this insight, the proposed LLM2TEA provides the LLM with the current genotype population and instructs it to reflect on them (as shown in Fig. 2(c)) prior to generating offspring. This reflection ensures the context-aware execution of the assigned mating operations, aligning genetic variation effectively with the objectives of the algorithm.

During offspring generation, the LLM is instructed to select parent genotypes from the current population and undergo crossover and/or mutation operations. Conventional EAs typically employ simple genetic crossover operators [32] or probabilistic models, such as Gaussian distributions [33], as mutation operators [34]. However, these traditional crossover and mutation techniques exhibit two main limitations that could inhibit the effectiveness of prompt search techniques.

Firstly, meaningful prompts in natural language often entail implicit constraints inherent in language sequences. Traditional crossover and mutation operators fail to account for these implicit constraints, potentially resulting in nonsensical prompts that may not effectively generate reasonable outputs. Secondly, considering that an LLM contains extensive prior knowledge, harnessing for crossover and mutation [35] can facilitate the generation of more coherent and contextually appropriate prompts by utilizing this embedded knowledge. Building on the above considerations, LLM plays the role of an evolutionary operator to facilitate the mating process. This process involves selecting parents, followed by crossover and/or mutation operations to construct a genotype, and finally, generating a design conditioned by the constructed genotype. In LLM2TEA, several LLM-based self-mating, same-domain mating, and cross-domain mating operators are proposed, which are described as follows:

- *Self-Mating*: Instruction is provided to the LLM to select a parent genotype from the current population and construct a new genotype based on its context.
- *Same-Domain Mating*: Instruction is provided to the LLM to randomly select two or more same-domain parent genotypes, constructing a new genotype by combining words from both. The LLM can mutate the new same-domain genotype further by randomly replacing some words.
- *Cross-Domain Mating*: Instruction is provided to the LLM to randomly select two or more parent genotypes from different domains by combining words from both. The LLM can mutate the new cross-domain genotype further by randomly replacing some words.

During the mating process, if both selected parents belong to the same domain, the same-domain mating operator (shown in Fig. 2 (e)) kicks in; otherwise, a cross-domain mating operator (shown in Fig. 2 (f)) takes effect. Additionally, the self-mating operator (shown in Fig. 2 (d)) involving a single parent is triggered with a certain probability. The resulting genotypes

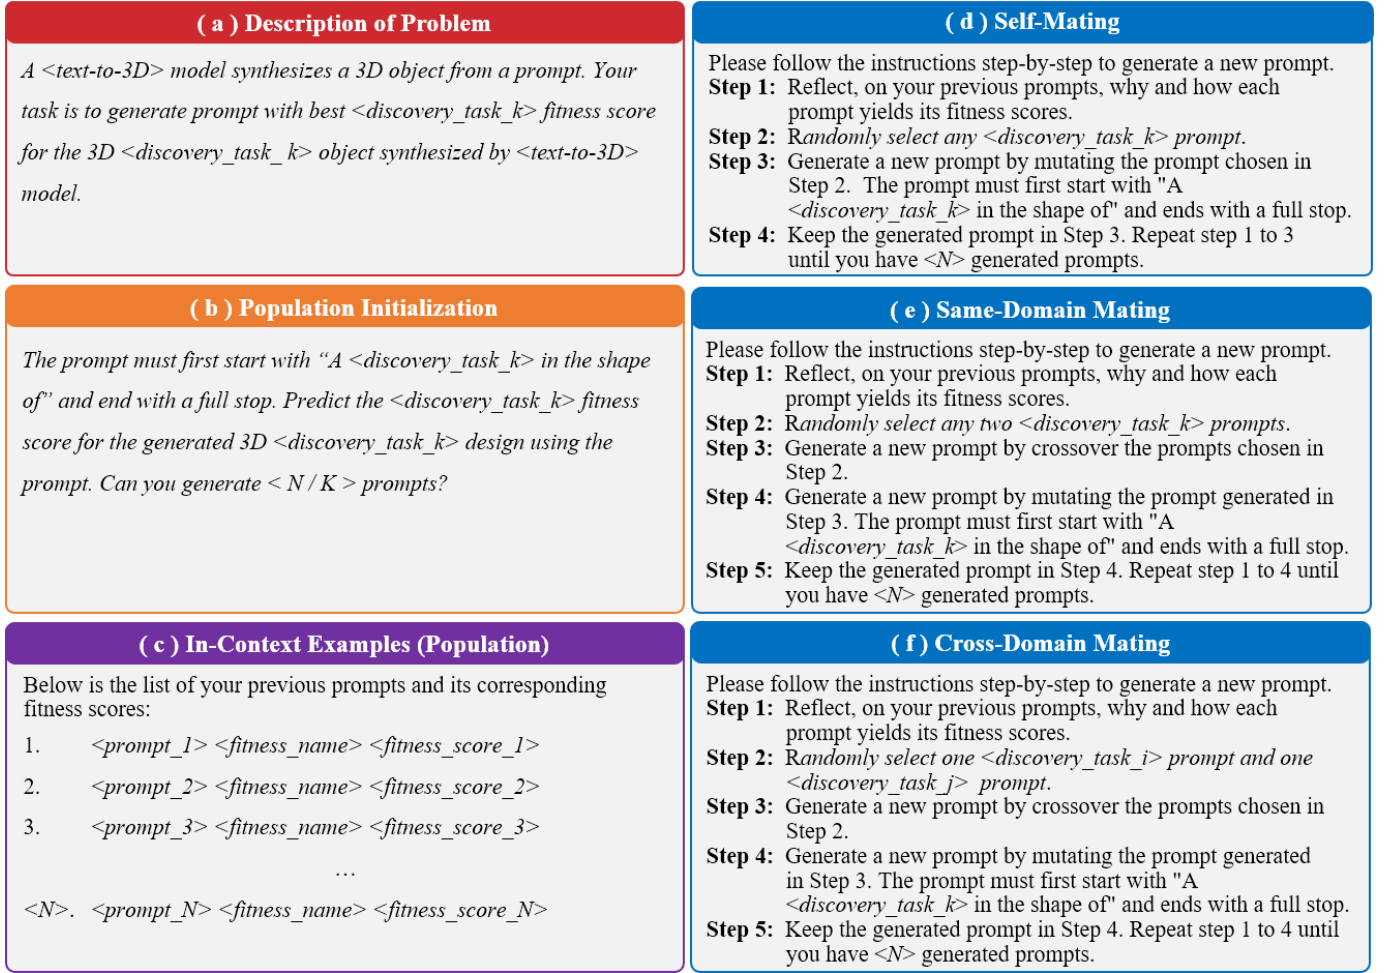


Fig. 2. Instruction sets for the LLM to function as an evolutionary optimizer. LLM2TEA enables the LLM to initialize and evolve prompts guided by various physical and visual criteria in the evolution process.

are passed into a text-to-3D generative model, thereby synthesizing new 3D designs. The generated 3D design and its associated prompt are the designs of an offspring genotype. All the offspring genotypes will form the offspring population O_t and are evaluated based on the physical and visual criteria. Subsequently, the offspring population is combined with the current population P_t as the joint population. Any duplicate genotypes found in this joint population will be removed in the current generation population, retaining the unique genotype originating from the offspring population. The resulting joint population undergoes selection pressure based on the surviving genotypes.

3) *Selection*: A scalar fitness measure to compare genotypes across domains is considered. Specifically, given K tasks and N genotypes in the joint population where $j \in \{1, \dots, N\}$, a set of factorial ranks $\{r_k^{(1)}, \dots, r_k^{(N)}\}$ for the K tasks is determined by finding the ordered position of every genotype with respect to all other genotypes' factorial cost. Here, this factorial cost is the genotype's fitness score. As such, each genotype in the joint population will have a set of factorial ranks for all tasks. Thereafter, the scalar fitness value φ_j for the j -th genotype is computed as $\min_{k \in \{1, \dots, K\}} \{r_k^{(j)}\}$. The scalar fitness values of the joint population are used in

any evolutionary selection technique to eliminate less fitting genotypes, allowing the surviving individuals to form the next generation of population in the evolutionary process.

IV. EXPERIMENTS SETUP

The experimental setup to evaluate various LLM2TEA capabilities and the discovered designs is presented here. First, the configurations for various test scenarios are introduced. Next, the experimental settings of the LLM2TEA are elaborated.

A. Test Scenarios

The performance of LLM2TEA is observed in the context of engineering design instances. Following Equations (1) and (2), the experiments are set up as a two-task design problem that targets to find novel car and airplane designs concurrently in a single EMT run. As such, the *car* domain is $\mathbf{t}_1 = \text{"A car in the picture"}$, while the *airplane* domain is $\mathbf{t}_2 = \text{"An airplane in the picture"}$. The effectiveness of LLM2TEA with ChatGPT 3.5 [36] in driving evolution and Shape-E [14] as the auxiliary text-to-3D generative model is considered. The search space $\mathbb{Z}^{L \times V}$ for effective prompts is thus $L = 77$ and $V \approx 50K$ for Shape-E, while a maximum

sequence context length of $C \approx 16K$ and $V_c \approx 50K$ for ChatGPT 3.5. In terms of f_{visual} in Equation (1) and (2), BLIP-2 VLM [22] is used as the evaluator to assess the visual performance of the generated designs. This visual criterion has a contribution weight of $\alpha = 0.55$, taking priority over the physical criteria in the objective. In addition, $c \geq 0.5$ and $\bar{c} \leq 1.0$ in Equation (2) are set to constrain the exploration of generated designs within and near domain boundaries. Four test scenarios are considered to study the effectiveness of the LLM2TEA:

- *Test Scenario I:* This test scenario assesses the design diversity attained by LLM2TEA in a multi-objective setting that aims to minimize aerodynamic drag in both car and airplane designs, simultaneously maximizing aerodynamic lift in airplanes while minimizing aerodynamic lift in car designs.
- *Test Scenario II:* This test scenario considers a minimization of the projected frontal area [37] as f_{phy} .
- *Test Scenario III:* This test scenario minimizes aerodynamic drag criterion.
- *Test Scenario IV:* This test scenario assesses the creativity capability of LLM2TEA in discovering novel designs. Specifically, the two-task optimization tasks are defined as $\mathbf{t}_1 = \text{“banana car”}$ and $\mathbf{t}_2 = \text{“banana airplane”}$. The generated designs are grounded in real-world performance-based optimization criteria, allowing for an assessment of creativity while ensuring conformity to real-world requirements.

Aerodynamic drag and lift of generated designs are computed using OpenFOAM [38] and normalized. OpenFOAM is an open-source computational fluid dynamics software toolkit primarily used for simulating fluid flow, heat transfer, turbulence, chemical reactions, and other related physical phenomena.

B. Parameter Settings

In all experiments, the population size is set to $N = 20$, and consistent random seeds are used throughout the evolution, with a terminating criterion of up to 20 generations. This results in up to 400 designs being generated and evaluated by the text-to-3D generative model in a single run. The classic tournament selection algorithm [39] is used to select the N genotypes that form the next generation of designs. Each experiment run is executed on a single shared compute node equipped with Intel Xeon Silver 64 CPU cores clocked at 2.10 GHz, 128 GB of RAM, and three Nvidia Quadro GV100 GPUs (each with 32 GB of RAM), with both evolution and simulation run in parallel.

C. Evaluating Designs

In the context of LLM2TEA, an important observation is its ability to generate designs that deviate from its conventional form. This capability is particularly evident in cross-domain design scenarios, where integrating functional principles and constraints from multiple domains often produces artifacts with fused structural or behavioral traits, resulting in novelty. However, in existing generative design research in literature [40]–[46], there is no clear definition or standard for

the concept of “novel” design. This paper defines design novelty by examining the distinction between “conventional” and “unconventional” shapes. Specifically, randomly sampling shapes from a generative model based on a known domain \mathbf{t} (e.g., a car or an airplane) is considered conventional, i.e., designs sampled as $\mathbf{x} \sim p_\theta(\mathbf{x}|\mathbf{t})$. Designs with shapes deviating from this conventional set are considered novel designs.

To quantify novelty, a Novelty Score metric is proposed based on this idea. Using a VLM, a generated 3D design is compared with the visual bias of a target concept specified in text form. Specifically, the VLM is tasked to measure the similarity distance between the design and the task label (such as “A car” or “A airplane”). This measurement can indicate whether the design is visually similar to a conventional concept of the vehicle task that VLM has already learned and whether the text-to-3D generative model can be synthesized. As such, *Novelty Score* metric is introduced which is formulated as:

$$NS(\mathbf{x}, \hat{\mathbf{z}}_k) = g_{vlm}(\mathbf{x}, \hat{\mathbf{z}}_k) - \mathbb{E}_{p_\theta(\hat{\mathbf{x}}|\hat{\mathbf{z}}_k)} [g_{vlm}(\hat{\mathbf{x}}, \hat{\mathbf{z}}_k)], \quad (3)$$

where $\hat{\mathbf{z}}_k$ is the k -th task label with the format “A < \mathbf{t}_k >” and $g_{vlm}(\cdot)$ is the independent VLM (e.g., CLIP [21]) to measure the similarity distance between a design and its corresponding task label. This VLM is used in the Shape-E model to train and generate 3D designs, aligning with the visual bias inherent in the text-to-3D generative model. Notice that when $NS(\cdot) \leq 0$, the design is considered to exhibit only conventional features. Conversely, when $NS(\cdot) > 0$, the design is considered to be “novel”, indicating a deviation from the conventional forms known to the VLM. On a separate note, Monte Carlo estimation is adopted here to approximate $\mathbb{E}_{p_\theta(\hat{\mathbf{x}}|\hat{\mathbf{z}}_k)} [g_{vlm}(\hat{\mathbf{x}}, \hat{\mathbf{z}}_k)]$ using 1,000 generated samples per task as the baseline mean, taking into account the variety of conventional shapes that resemble the target task.

V. RESULTS AND DISCUSSIONS

This section presents the experiment results in the following order. First, the designs discovered during the experiments are showcased. Next, the diversity of these designs is analyzed, followed by a discussion of the various components of the LLM2TEA to evaluate its capabilities and effectiveness. Finally, experiments demonstrating the 3D printing of the generated designs are presented.

A. Results of Discovered Designs

Some of the most representative examples from all test scenarios of vehicle designs uncovered in the text-to-3D generative model, Shape-E, are presented in Fig. 3. Vehicle designs from the baseline exhibit conventional vehicular features, as shown in Fig. 3(a) and (b), and have a distinct contrast with the novel designs found with the LLM2TEA, as illustrated in Fig. 3(c) and (d). Such novel designs that deviate from conventional form result from guiding the generative models to synthesize complex prompts. These prompts lead to designs that not only satisfy the physical criterion but also fulfill the visual specification of a car or an airplane. A clear distinction can be observed between designs generated with LLM2TEA and those from the baseline model. The




Fig. 3. Examples of designs generated by text-to-3D generative model are shown. In contrast to the vehicle body shapes generated directly by the model without the LLM2TEA (Fig. 3(a) and (b)), the designs discovered with LLM2TEA deviate from these conventional shapes, since the model are guided to synthesize designs that not only satisfy physical criteria but also embody the visual characteristics of a car or an airplane, as illustrated by some of the representative example designs in Fig. 3(c) and (d). The cross-domain search enabled by the multitask setting facilitates the discovery of hybrid vehicle designs, as most evident in Fig. 3(d).

cross-domain knowledge learning capability of LLM2TEA enables the genetic material from parent prompts across different domains to undergo crossover and mutation, thereby facilitating the transfer of knowledge across domains. This process generates offspring prompts that drive the discovery of hybrid vehicle designs. These designs feature unconventional body shapes and structures, such as banana-like curvatures or extended rear spoilers, challenging traditional car and airplane design concepts.

Remark: In LLM2TEA, the implicit genetic transfer occurring during cross-domain mating facilitates the creation of hybrid designs that combine the distinctive traits of each parent genotype from different domains. On a separate note, LLM and text-to-X generative models known for its hallucination phenomenon [47] may lead to synthesizing outputs that deviate from the factual or the input prompt. Such a hallucination outcome is not always detrimental but can foster positive creativity [48]. Both of these algorithmic properties create the opportunity for emerging design trends, facilitating the discovery of novel designs.

B. Diversity of Designs

In Test Scenario I, the best-performing designs from the last five generations in a multi-objective setting of a single run using LLM2TEA are compared against baseline designs generated with a text-to-3D generative model without LLM2TEA. As shown in Fig. 4, the designs discovered with the LLM2TEA exhibit superior aerodynamic performance compared to the

TABLE I
COMPARISON OF HYPERVOLUME BETWEEN BASELINE AND LLM2TEA IN SINGLE EVOLUTIONARY RUN

Method	Airplane Task	Car Task
Baseline	0.3111	0.4907
LLM2TEA	0.8515	0.9702

baseline. Specifically, airplane designs demonstrated better lift (\tilde{C}_l) and drag (\tilde{C}_d) than baseline designs, with similar trends observed for car designs. Apart from that, the baseline text-to-3D generative model tends to discover designs clustered in localized regions, whereas LLM2TEA produces more diverse designs, as evident from the visual inspection of the generated airplane designs. To quantitatively assess the diversity, Table I presents the hypervolume indicator [49] measured over the physics objectives (\tilde{C}_l and \tilde{C}_d) space, which ranges from 0.0 to 1.0. LLM2TEA achieves hypervolume values 2.74 greater for airplane designs and 1.97 times for car designs compared to the baseline. These results demonstrated that the proposed LLM2TEA effectively explores design spaces to uncover novel designs with strong physics performance.

C. Effectiveness of LLM2TEA

A series of experiments was conducted to objectively examine the capabilities of LLM2TEA to discover novel solutions. The findings are presented here.

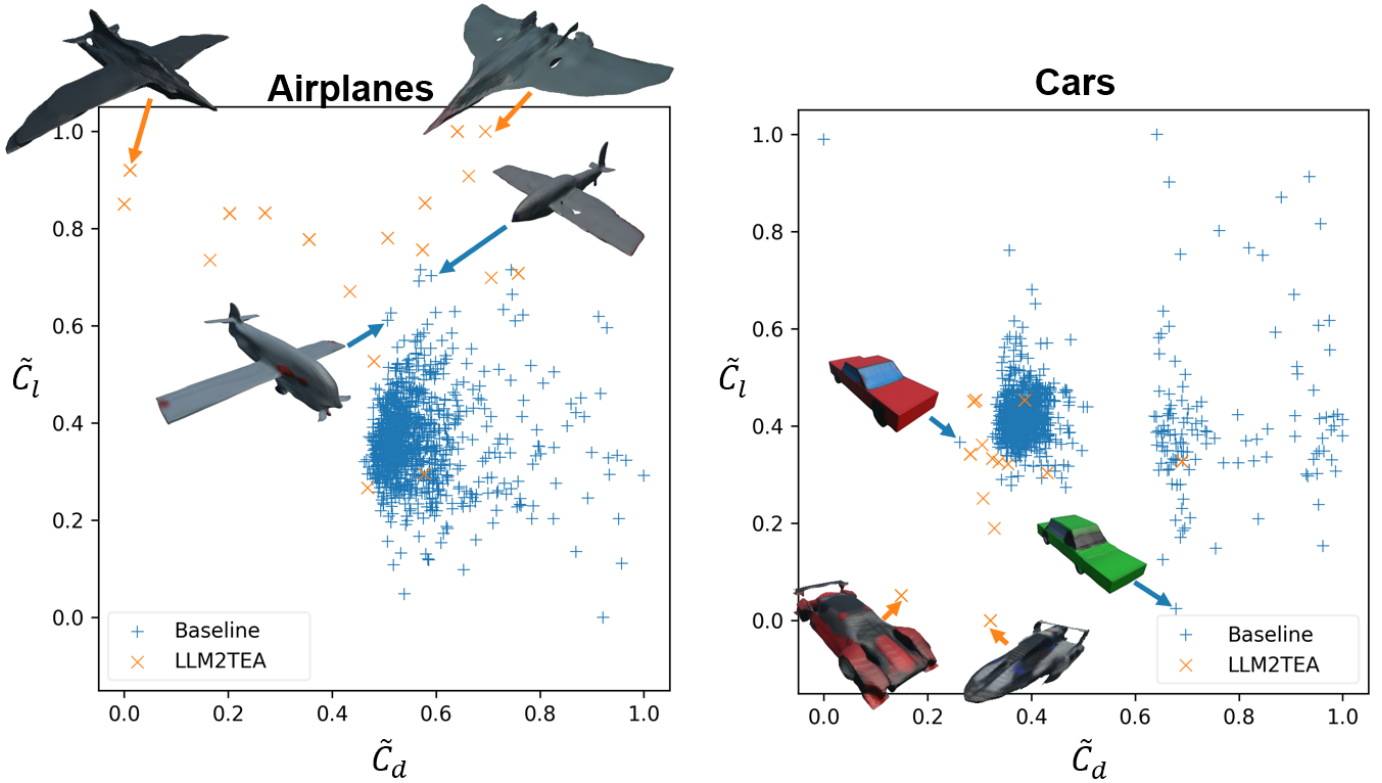


Fig. 4. Test Scenario I compares between baseline designs generated by the baseline text-to-3D generative model without LLM2TEA to designs discovered with LLM2TEA during the last five generations of a single evolutionary run. The designs discovered with LLM2TEA demonstrated superior expected aerodynamic performance (exhibiting higher normalized lift (\tilde{C}_l) and lower normalized drag (\tilde{C}_d) scores for airplanes, as well as lower \tilde{C}_l and \tilde{C}_d scores for cars) and are visually more novel than the baseline designs. (Note that the illustrations presented here are not drawn to scale.)


1) *Effect of Evaluation Criteria:* This subsection examines the effects of the objective functions defined in Section III-A, demonstrating their ability to guide LLM2TEA toward finding the best-performing designs. As shown in Fig. 5, when guided solely by visual criteria, LLM2TEA generates vehicle designs with wide body shapes and high roofs that retain the characteristics of a car or airplane but exhibit poor aerodynamic performance (Fig. 5(a)). In contrast, designs evaluated solely on physical criteria will result in the production of more highly fragmented designs, such as the example presented in Fig. 5(b). Therefore, incorporating both physical and visual criteria enables LLM2TEA to discover practical designs with body shapes that achieve the lowest drag (Fig. 5(d)) compared to those in Fig. 5(a) through (c). Notably, over 73% of the generated vehicle designs outperform the top 1% designs produced by the text-to-3D baseline in terms of physical performance. These results suggest that LLM2TEA effectively uncovers novel yet practical designs.

2) *Evolution Trends of LLM2TEA:* Table II presents LLM2TEA average fitness scores across generations in the single evolutionary run. A distinct trend of gradual improvements in the initial phases can be observed, particularly for the airplane design task. This progression becomes more pronounced in the latter half of the run, indicating that the evolutionary search led to the discovery of better aerodynamic designs. Further examination of the results in all test scenarios supports the observed trend (Figs. 6, 7, and 8). While the initial


vehicle designs exhibit body shapes resembling conventional forms, LLM2TEA enables the designs to evolve and deviate from their conventional forms, with more significant changes observed from the mid-generation phase of the evolutionary search.

3) *Behavior of Mating Operations:* In Fig. 7, notice the less frequently used words included in the car prompts, such as “wings” and “jet”. Such prompt constructs are the result of the LLM2TEA mating operations. Self-mating and same-domain mating operations construct same-domain prompts that utilize commonly used terms, resulting in vehicle designs with conventional body shapes. In contrast, cross-domain mating operation, with its cross-concept knowledge integration, enables the crossbreeding between different domains by combining shape-defining words from these domains to produce a new genotype. The resulting phenotype generated with this new genotype exhibits unconventional body shapes with a hybrid appearance, combining distinctive features of these domains while still satisfying the physical criteria (Fig. 9).

Since the LLM2TEA cross-concept knowledge transfer mechanism operates by transferring representative words from one task to another, the effectiveness of this capability can be observed by examining the percentage overlap between each task’s vocabulary set. The vocabulary set for each task comprises representative word tokens of all the prompts generated throughout the evolutionary run. For a fair comparison, frequently occurring stopwords (such as “a”, “the”, “is”, etc)



“A airplane in the shape of a mysterious shadow.”




“A car in the shape of a compact city hatchback with a practical design and efficient fuel consumption.”

(a) LLM2TEA with Visual Criteria Only



“A airplane in the shape of a slender arrow piercing gracefully through the fog.”




“A car in the shape of branched chain.”

(b) LLM2TEA with Physical Criteria Only



“A airplane in the shape of a modern cutting-edge canard stealth jet blending state-of-the-art luxury technology with advanced high-speed fuselage elements for exceptional drag reduction and efficiency.”




“A car in the shape of a modern compact sedan with a sleek profile and integrated LED headlights.”

(c) LLM2TEA with Physical & Visual Criteria



Predefined Prompt:

“An airplane”



Predefined Prompt:

“A car”

(d) Baseline

Fig. 5. Examples of text prompts (right) used to generate the designs (left) are shown, with colors indicating the magnitude of the fluid velocity, where brighter colors correspond to higher velocity. When LLM2TEA evaluates using only visual criteria, the resulting vehicles tend to have rectangular shapes and high roofs, which correspond to poor drag performance, as illustrated in Fig. 5 (a). Conversely, relying solely on the physical criterion will degrade the search performance with more highly fragmented vehicle designs discovered, such as those shown in Fig. 5 (b). As such, by considering both physical and visual criteria simultaneously, LLM2TEA can discover practical designs (Fig. 5 (c)) featuring body shapes that yield better drag performance compared to baseline (Fig. 5 (d)).

TABLE II
POPULATION MEAN FITNESS SCORE AND ITS VARIANCE ACROSS GENERATIONS IN SINGLE EVOLUTIONARY RUN

Task	Initial	Gen. 5	Gen. 10	Gen. 15	(Final) Gen. 20
Airplane	0.2683 ± 0.0012	0.1596 ± 0.0015	0.1485 ± 0.0014	0.0662 ± 0.0415	-0.0216 ± 0.0427
Car	0.0981 ± 0.0001	0.0609 ± 0.0001	-0.0531 ± 0.0514	-0.0336 ± 0.0343	-0.1339 ± 0.1388

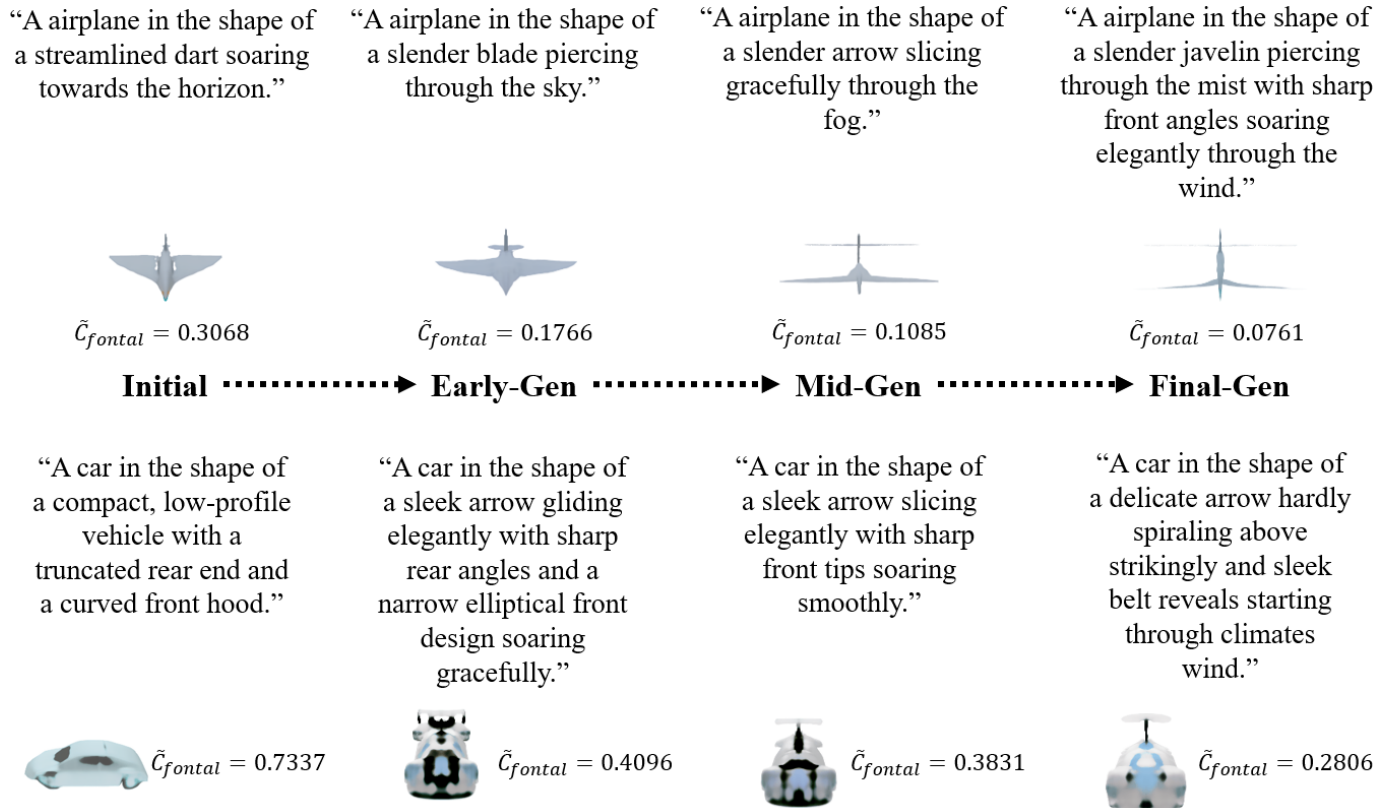


Fig. 6. Example of the evolution of prompts and their corresponding vehicle designs during an LLM2TEA evolutionary run in Test Scenario II, which aims to minimize the projected frontal area. Note that the technical indicator ($\tilde{C}_{frontal}$) provided is a computed value from the evaluator and does not reflect the actual real-world geometry.

are pruned from the vocabulary set. Furthermore, all word tokens are reduced to its base word form by transforming the morphological variants to the base word. The resultant vocabulary set comprises the word tokens most representative of each task. Henceforth, an 38.38% overlap in the vocabulary sets between tasks, indicating a proportion of word tokens shared between tasks. The results demonstrated the effectiveness of the cross-domain knowledge-driven approach, which transfers a significant amount of words between tasks with LLM2TEA.

4) *Efficiency and Scalability of LLM2TEA*: From our observations of the computational time of each key component in LLM2TEA, the computation time of the operators is negligible as proprietary LLMs is used, which runs on third-party servers. However, 72.7% of one generation total time is used in the OpenFoam physics simulation, which is typical in engineering design optimization problems. The remaining time is used for visual evaluations, accounting for 20.5% of the time, and text-to-3D generation utilizes the remaining 6.8% of the time.

Similar to existing engineering design methods, where scaling up for bigger problems requires parallelizing physical simulation computations, LLM2TEA efficiently uses multicore CPU processing to compute the physical performance of generated designs. However, in contrast to existing evolutionary methods that require a large population size to yield good results, LLM2TEA can achieve good results with a small population size. Additionally, since the generated designs are available in digital 3D formats, the designs can be scaled down to fit into a unit cube, further reducing the simulation computation requirements.

5) *Quantification of Novel Solutions*: The effectiveness of LLM2TEA in discovering novel solutions is examined from a visual perspective using the proposed Innovation Score metric defined in Equation (3). Results indicate that 69.12% of airplane designs and 46.23% of car designs generated by LLM2TEA exceed the computed mean baseline, underscoring the efficacy of the multitask evolutionary process in guiding

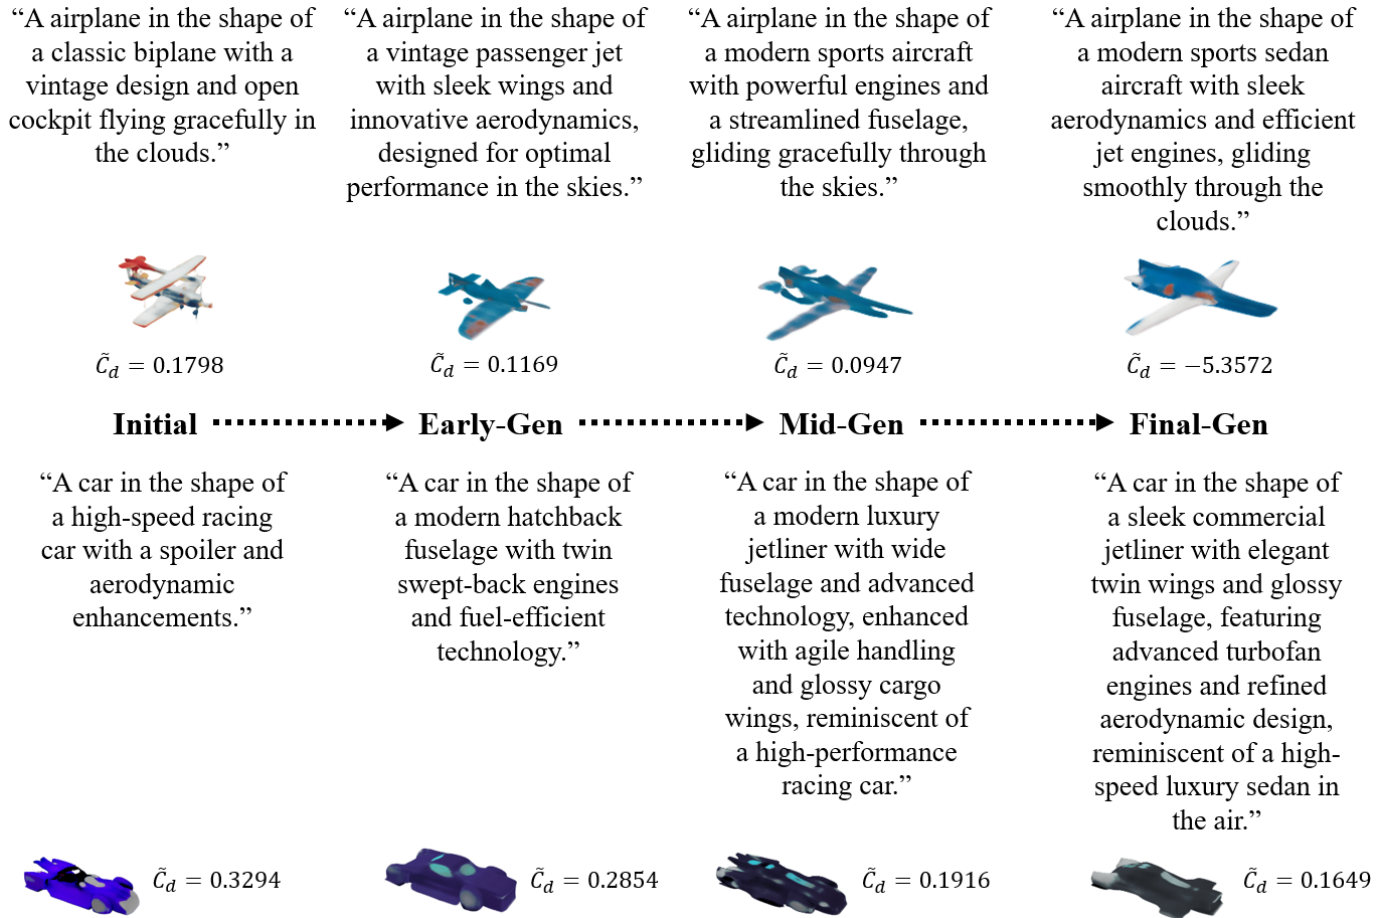


Fig. 7. Example of the evolution of the prompts and its corresponding vehicle designs during an LLM2TEA evolutionary run in Test Scenario III, which aims to minimize drag. Note that the technical indicator (\tilde{C}_d) provided is a computed value from computational physics simulation and does not reflect the true real-world performance.

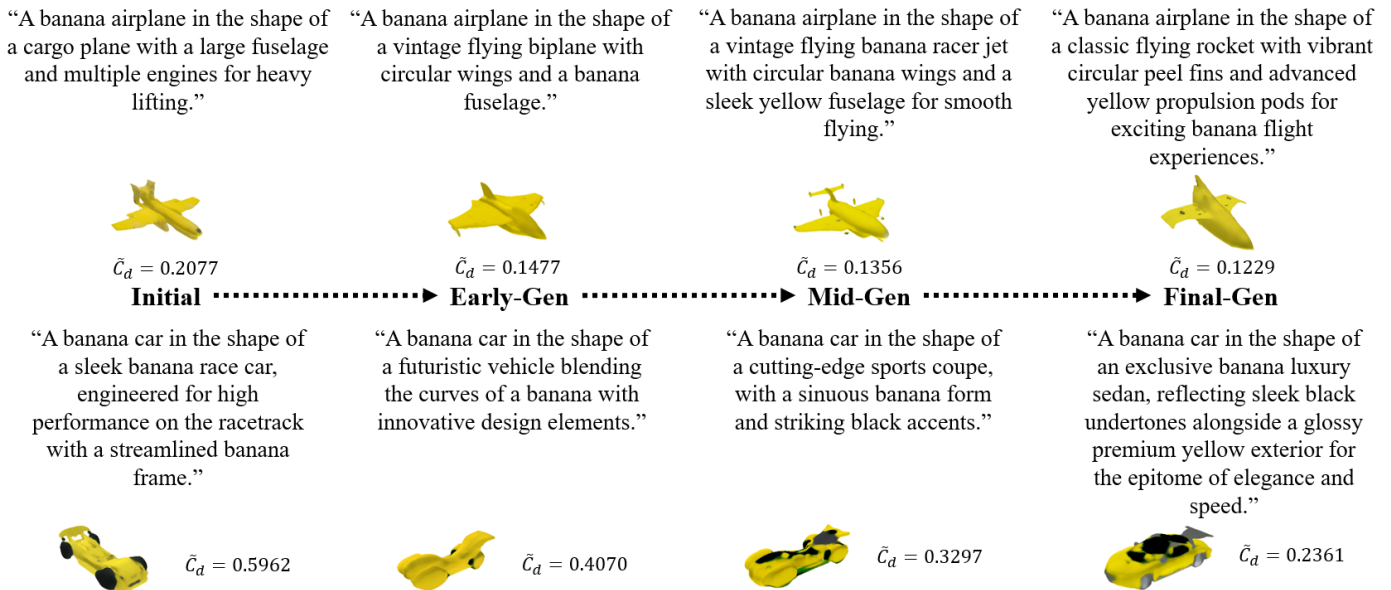


Fig. 8. Example of the evolution of the creative prompts and its corresponding creative vehicle designs during an LLM2TEA evolutionary run in Test Scenario IV, which aims to minimize the drag. Note that the technical indicator (\tilde{C}_d) provided is a computed value from computational physics simulation and does not reflect the true real-world performance.

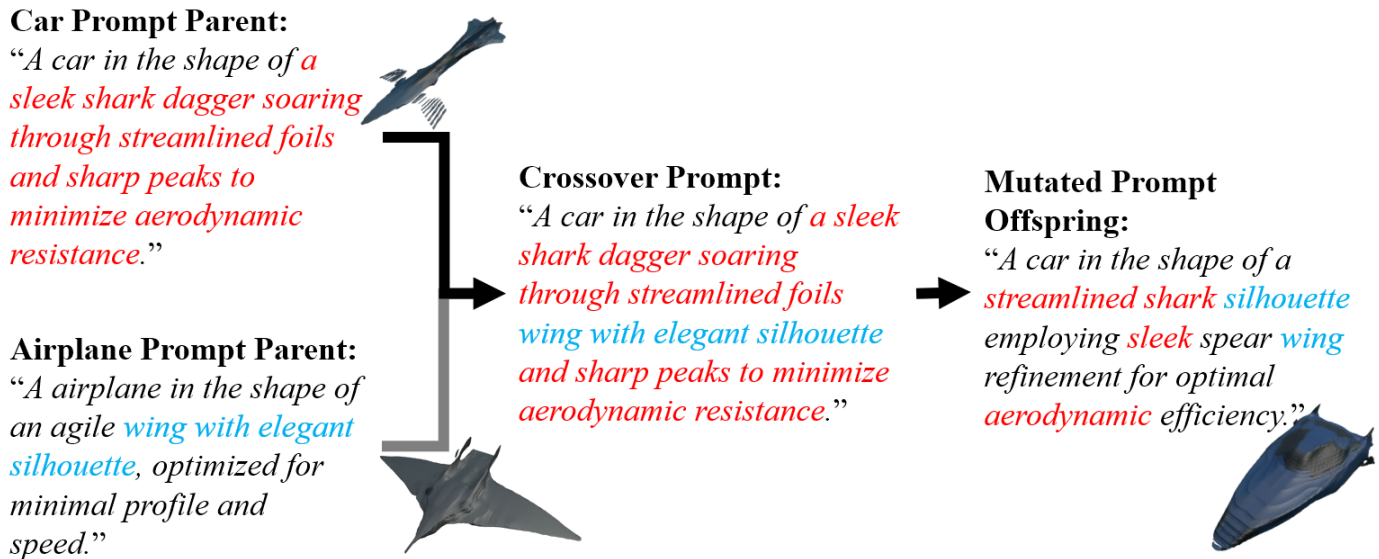


Fig. 9. Illustration of prompt crossbreeding between **car** and **airplane** domains, demonstrating multitask adaptation in the evolution of LLM2TEA.

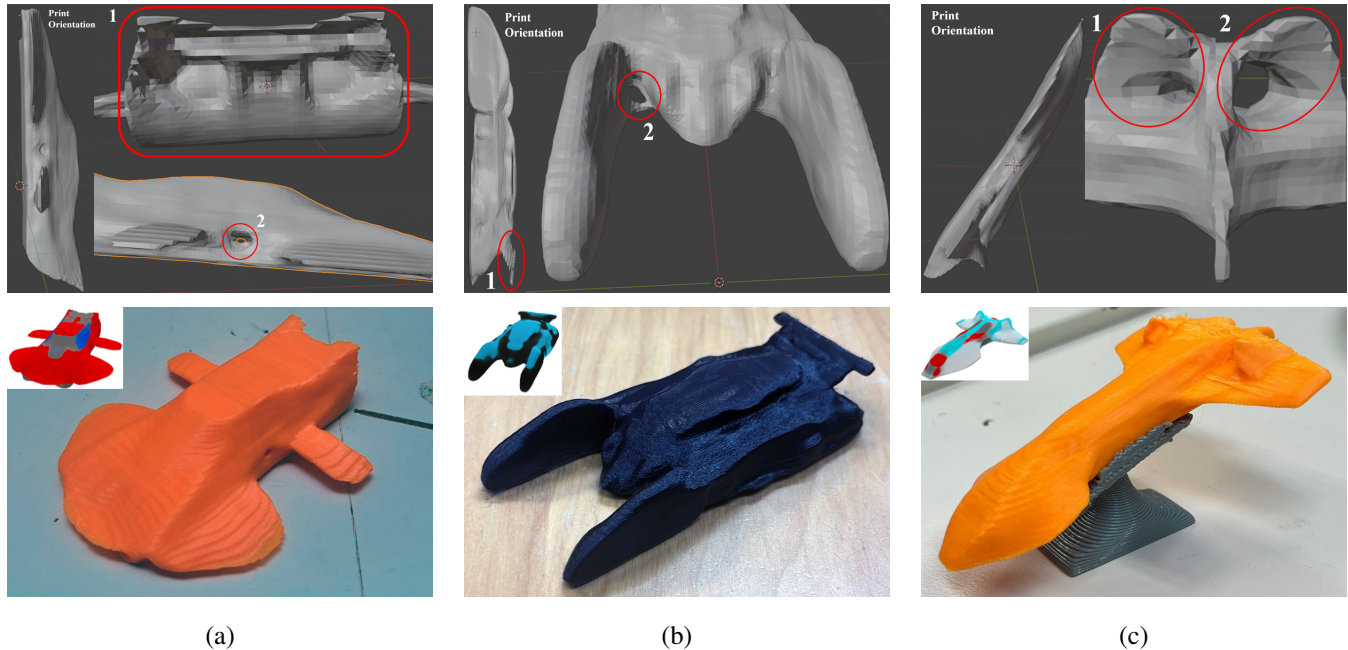


Fig. 10. Study conducted on 3D printability of discovered designs found that the majority of designs generated by text-to-3D generative model carry some degree of defects, as illustrated in the top row of this figure. Some defects are easily repairable without losing the novelty of the design shapes by thickening the weak or thin structures, filling holes, and strengthening the internal cavity. Furthermore, the vehicle design’s non-flat surfaces require printing orientation in a vertical position with various printing supports. The resulting 3D-printed designs, as shown in the bottom row, match closely with the original designs. This study demonstrated the potential use of GenAI for discovery and fast design prototyping, opening up new avenues for research and innovation.

generative models toward the discovery of cross-domain novel designs. This further validates LLM2TEA’s capability to yield unconventional solutions not only across domains but also within a single domain, as illustrated in Fig. 3(c) and (d), where the resulting designs exhibit more streamlined, futuristic, and sleeker body shapes that deviate from conventional norms.

D. 3D Design Printing

The designs discovered with LLM2TEA can be 3D printed. Since the text-to-3D generative model (Shape-E) tends to

generate designs that are non-conforming to manufacture in real-world conditions, most designs uncovered in the text-to-3D generative model possessed some degree of defects and fragmentation. For instance, as shown in the top row of Fig. 10, common defects, such as rough surfaces, sharp edges, open holes, and weak or thin structures, especially within internal cavities, can be easily resolved with minor repairs and infills to the designs before 3D printing. On the other hand, the non-flat and intricate surface details require printing orientation in a vertical position with minimal breakaway or tree-like printing supports. Additionally, the 3D designs require appropriate

scaling to preserve its detailed geometric features during 3D printing.

Taking these factors into account, off-the-shelf SLA 3D printing is used due to its ability to retain such intricate details and its suitability for small-quantity prototyping. The resulting 3D-printed designs are shown in the bottom row of Fig. 10, closely matching the digital visualization of the designs. The conceptual design process can be accomplished within weeks, from discovery to 3D printing the designs, with the start of examining the designs during a single evolutionary run in a matter of hours. This demonstration highlights the capability of LLM2TEA to transform AI-generated design outputs into tangible physical designs, underscoring its potential applications in fast design prototyping.

Concluding Remark: *We hope that the insights presented in this experiment discussion section will provide researchers a better understanding of the challenges in the problem of interest and address the identified technical generation issues in LLMs and GenAI models using our proposed strategies. There are several ways practitioners can set up LLM2TEA to suit their specific needs, including defining the task domains with the variable t_k , choosing a preferred physics simulator, and setting the hard constraint bound values \underline{c} and \bar{c} to control the exploration behavior. We learned that it is crucial to identify task domains that are sufficiently general, exhibit positive transfer of geometric characteristics between tasks, and integrate well with automatic mesh generation and simulation. Moreover, generated cross-domain designs should maintain validity for evaluations as exploration goes deeper into out-of-domain regions. In summary, the proposed method applies to any applications in science and engineering as long as the generated designs can be evaluated. These insights hold broader relevance to research communities, such as those in product design or animation, and other fields seeking to leverage LLM2TEA, LLMs, and GenAI models for novel design discovery.*

VI. CONCLUSION AND OUTLOOK

This paper presents LLM2TEA, the first agentic AI designer operating with GEM, leveraging LLMs to explore and discover novel, physically viable designs across multiple domains. By integrating LLM-guided evolutionary search, text-to-3D generation, semantic evaluation, and physics simulation, LLM2TEA empowers designers to navigate vast, cross-domain design spaces, unlocking novel solutions that ground aesthetic novelty with real-world practicality. In particular, LLM2TEA is an agentic AI designer capable of expanding the boundaries of conventional design by leveraging knowledge across multiple domains. It explores a broader design space that extends beyond individual areas and embraces their overlap through the paradigm of evolutionary multitasking. This approach promotes the discovery of novel cross-domain designs through effective implicit knowledge transfer. Experimental results demonstrate significant improvements in design diversity (up to 174%) and physical performance (with over 73% surpassing leading baselines). The successful 3D printing of several generated designs further validates the real-world feasibility

of LLM2TEA, underscoring its transformative potential for AI-driven, cross-domain innovation and practical design optimization.

While the empirical studies of the prototypical design optimization problem are an abstraction of real-world engineering design scenarios, the potential of the LLM2TEA in accelerating the conceptual design process and enhancing the accessibility of such solutions across multiple disciplines has been demonstrated. Nevertheless, this research direction remains in its infancy stage, and further explorations are encouraged in other scientific and engineering scenarios.

ACKNOWLEDGMENT

Melvin Wong gratefully acknowledges the financial support from Honda Research Institute Europe (HRI-EU). The work is also partially supported by the National Research Foundation (NRF), Singapore, through the AI Singapore Programme under the project titled ‘‘AI-based Urban Cooling Technology Development’’ (Award No. AISG3-TC-2024-014-SGKR), as well as by the College of Computing and Data Science (CCDS) at Nanyang Technological University (NTU). The authors would like to thank Mr. Jason Fow Kow, Mr. Lim Rui Yi Ray, Ms. Sun Sitong, and Ms. Swaminathan Navitraa for their assistance in 3D printing the AI-generated designs and supporting video production. Grammarly and ChatGPT were used as writing assistance tools to enhance the readability of the manuscript.

REFERENCES

- [1] M. Wong, Y. S. Ong, A. Gupta, K. K. Bali, and C. Chen, ‘‘Prompt evolution for generative ai: A classifier-guided approach,’’ in *2023 IEEE Conference on Artificial Intelligence (CAI)*. IEEE Computer Society, 2023, pp. 226–229. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/CAI54212.2023.00105>
- [2] A. Y. Radwan, K. M. Alasmari, O. A. Abdulbagi, and E. A. Alghamdi, ‘‘Sard: A human-ai collaborative story generation,’’ in *International Conference on Human-Computer Interaction*. Springer, 2024, pp. 94–105.
- [3] D. Podell, Z. English, K. Lacey, A. Blattmann, T. Dockhorn, J. Müller, J. Penna, and R. Rombach, ‘‘Sdxl: Improving latent diffusion models for high-resolution image synthesis,’’ *arXiv preprint arXiv:2307.01952*, 2023.
- [4] R. Huang, J. Huang, D. Yang, Y. Ren, L. Liu, M. Li, Z. Ye, J. Liu, X. Yin, and Z. Zhao, ‘‘Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models,’’ in *International Conference on Machine Learning*. PMLR, 2023, pp. 13 916–13 932.
- [5] ‘‘Veo 2,’’ Dec 2024. [Online]. Available: <https://deepmind.google/technologies/veo/veo-2/>
- [6] T. Rios, S. Menzel, and B. Sendhoff, ‘‘Large language and text-to-3d models for engineering design optimization,’’ in *2023 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2023, pp. 1704–1711.
- [7] M. Wong, T. Rios, S. Menzel, and Y. S. Ong, ‘‘Prompt evolutionary design optimization with generative shape and vision-language models,’’ in *2024 IEEE Congress on Evolutionary Computation (CEC)*, 2024, pp. 1–8.
- [8] J. Har and K. Tamma, *Advances in computational dynamics of particles, materials and structures*. John Wiley & Sons, 2012.
- [9] L. Regenwetter, A. H. Nobari, and F. Ahmed, ‘‘Deep generative models in engineering design: A review,’’ *Journal of Mechanical Design*, vol. 144, no. 7, p. 071704, 2022.
- [10] J. R. Martins and A. Ning, *Engineering design optimization*. Cambridge University Press, 2021.
- [11] Q. Xu, J. Liu, M. Wong, G. Jin, R. Lau, Y. S. Ong, S. Menzel, T. Rios, J.-H. Lim, and C. C. Ooi, ‘‘Towards end-to-end prompt-vision-physics neural network for fast design discovery,’’ in *2024 IEEE Conference on Artificial Intelligence (CAI)*. IEEE, 2024, pp. 1432–1435.

- [12] Q. Xu, J. Liu, M. Wong, C. Chen, and Y. S. Ong, "Looks great, functions better: Physics compliance text-to-3d shape generation," in *International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2025.
- [13] A. Nichol, H. Jun, P. Dhariwal, P. Mishkin, and M. Chen, "Point-e: A system for generating 3d point clouds from complex prompts," 2022.
- [14] H. Jun and A. Nichol, "Shap-E: Generating Conditional 3D Implicit Functions," *arXiv preprint arXiv:2305.02463*, 2023. [Online]. Available: <https://arxiv.org/abs/2305.02463>
- [15] M. Cubes, "Construction of topologically correct isosurface," 1995.
- [16] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. H. Chi, Q. V. Le, and D. Zhou, "Chain-of-thought prompting elicits reasoning in large language models," in *Proceedings of the 36th International Conference on Neural Information Processing Systems*, ser. NIPS '22. Red Hook, NY, USA: Curran Associates Inc., 2022.
- [17] S. Yao, D. Yu, J. Zhao, I. Shafraan, T. L. Griffiths, Y. Cao, and K. Narasimhan, "Tree of thoughts: deliberate problem solving with large language models," in *Proceedings of the 37th International Conference on Neural Information Processing Systems*, ser. NIPS '23. Red Hook, NY, USA: Curran Associates Inc., 2023.
- [18] S. Diao, Z. Huang, R. Xu, X. Li, Y. Lin, X. Zhou, and T. Zhang, "Black-box prompt learning for pre-trained language models," *arXiv preprint arXiv:2201.08531*, 2022.
- [19] R. Pryzant, D. Iter, J. Li, Y. T. Lee, C. Zhu, and M. Zeng, "Automatic prompt optimization with" gradient descent" and beam search," *arXiv preprint arXiv:2305.03495*, 2023.
- [20] Y. Wen, N. Jain, J. Kirchenbauer, M. Goldblum, J. Geiping, and T. Goldstein, "Hard prompts made easy: gradient-based discrete optimization for prompt tuning and discovery," in *Proceedings of the 37th International Conference on Neural Information Processing Systems*, ser. NIPS '23. Red Hook, NY, USA: Curran Associates Inc., 2023.
- [21] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning transferable visual models from natural language supervision," in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 8748–8763. [Online]. Available: <https://proceedings.mlr.press/v139/radford21a.html>
- [22] J. Li, D. Li, S. Savarese, and S. Hoi, "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models," *arXiv preprint arXiv:2301.12597*, 2023.
- [23] C. Yang, X. Wang, Y. Lu, H. Liu, Q. V. Le, D. Zhou, and X. Chen, "Large language models as optimizers," *arXiv preprint arXiv:2309.03409*, 2023.
- [24] Q. Guo, R. Wang, J. Guo, B. Li, K. Song, X. Tan, G. Liu, J. Bian, and Y. Yang, "Connecting large language models with evolutionary algorithms yields powerful prompt optimizers," *arXiv preprint arXiv:2309.08532*, 2023.
- [25] A. Gupta, Y. S. Ong, and L. Feng, "Insights on transfer optimization: Because experience is the best teacher," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 1, pp. 51–64, 2018.
- [26] H. Liu, F. Zhao, L. Wang, T. Xu, and C. Dong, "Evolutionary multitasking memetic algorithm for distributed hybrid flow-shop scheduling problem with deterioration effect," *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 1390–1404, 2025.
- [27] C. Yang, Q. Chen, Z. Zhu, Z.-A. Huang, S. Lan, and L. Zhu, "Evolutionary multitasking for costly task offloading in mobile-edge computing networks," *IEEE Transactions on Evolutionary Computation*, vol. 28, no. 2, pp. 338–352, 2024.
- [28] T. Rios, B. van Stein, T. Bäck, B. Sendhoff, and S. Menzel, "Multitask shape optimization using a 3-d point cloud autoencoder as unified representation," *IEEE Transactions on Evolutionary Computation*, vol. 26, no. 2, pp. 206–217, 2022.
- [29] L. Feng, Y. Huang, L. Zhou, J. Zhong, A. Gupta, K. Tang, and K. C. Tan, "Explicit evolutionary multitasking for combinatorial optimization: A case study on capacitated vehicle routing problem," *IEEE transactions on cybernetics*, vol. 51, no. 6, pp. 3143–3156, 2020.
- [30] S. Liu, C. Chen, X. Qu, K. Tang, and Y. S. Ong, "Large language models as evolutionary optimizers," *arXiv preprint arXiv:2310.19046*, 2023.
- [31] F. Liu, X. Lin, Z. Wang, S. Yao, X. Tong, M. Yuan, and Q. Zhang, "Large language model for multi-objective evolutionary optimization," *arXiv preprint arXiv:2310.12541*, 2023.
- [32] G. Pavai and T. Geetha, "A survey on crossover operators," *ACM Computing Surveys (CSUR)*, vol. 49, no. 4, pp. 1–43, 2017.
- [33] I. Ono, "A real-coded genetic algorithms for function optimization using unimodal normal distribution crossover," in *Proc. 7th Int. Conf. on Genetic Algorithms, 1997*, 1997, pp. 246–253.
- [34] R. Hinterding, "Gaussian mutation and self-adaption for numeric genetic algorithms," in *Proceedings of 1995 IEEE International Conference on Evolutionary Computation*, vol. 1. IEEE, 1995, p. 384.
- [35] E. Meyerson, M. J. Nelson, H. Bradley, A. Gaier, A. Moradi, A. K. Hoover, and J. Lehman, "Language model crossover: Variation through few-shot prompting," *arXiv preprint arXiv:2302.12170*, 2023.
- [36] OpenAI, "Chatgpt 3.5." [Online]. Available: <https://platform.openai.com/docs/api-reference>
- [37] M. de Berg, *Computational geometry: algorithms and applications*. Springer Science & Business Media, 2000.
- [38] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby, "A tensorial approach to computational continuum mechanics using object-oriented techniques," *Computers in physics*, vol. 12, no. 6, pp. 620–631, 1998.
- [39] B. L. Miller, D. E. Goldberg *et al.*, "Genetic algorithms, tournament selection, and the effects of noise," *Complex systems*, vol. 9, no. 3, pp. 193–212, 1995.
- [40] H. O. Demirel, M. H. Goldstein, X. Li, and Z. Sha, "Human-centered generative design framework: an early design framework to support concept creation and evaluation," *International Journal of Human-Computer Interaction*, vol. 40, no. 4, pp. 933–944, 2024.
- [41] A. Di Filippo, M. Lombardi, A. Lorusso, F. Marongiu, and D. Santaniello, "Generative design for project optimization (s)," in *DMSVIVA*, 2021, pp. 110–115.
- [42] M. Martorelli and A. Gloria, "Strategies and generative design towards the development of innovative products," in *Springer Handbook of Additive Manufacturing*. Springer, 2023, pp. 269–286.
- [43] L. S. Furtado, J. B. Soares, and V. Furtado, "A task-oriented framework for generative ai in design," *Journal of Creativity*, vol. 34, no. 2, p. 100086, 2024.
- [44] R. Jaisawal and V. Agrawal, "Generative design method (gdm) – a state of art," *IOP Conference Series: Materials Science and Engineering*, vol. 1104, no. 1, p. 012036, mar 2021. [Online]. Available: <https://dx.doi.org/10.1088/1757-899X/1104/1/012036>
- [45] J. Mountstephens and J. Teo, "Progress and challenges in generative product design: A review of systems," *Computers*, vol. 9, no. 4, p. 80, 2020.
- [46] A. Agkathidis, "Generative design methods," in *Proceedings of eCAADe*, 2015, pp. 47–55.
- [47] N. Maleki, B. Padmanabhan, and K. Dutta, "Ai hallucinations: a misnomer worth clarifying," in *2024 IEEE Conference on Artificial Intelligence (CAI)*. IEEE, 2024, pp. 133–138.
- [48] X. Jiang, Y. Tian, F. Hua, C. Xu, Y. Wang, and J. Guo, "A survey on large language model hallucination via a creativity perspective," 2024.
- [49] H. Ishibuchi, R. Imada, Y. Setoguchi, and Y. Nojima, "How to specify a reference point in hypervolume calculation for fair performance comparison," *Evolutionary computation*, vol. 26, no. 3, pp. 411–440, 2018.