

Democratizing Energy Management with LLM-Assisted Optimization Autoformalism

Ming Jin, Bilgehan Sel, Fnu Hardeep, Wotao Yin

Abstract—This paper introduces a method for personalizing energy optimization using large language models (LLMs) combined with an optimization solver. This approach, termed *human-guided optimization autoformalism*, translates natural language specifications into optimization problems, enabling LLMs to handle various user-specific energy-related tasks. It allows for nuanced understanding and nonlinear reasoning tailored to individual preferences. The research covers common energy sector tasks like electric vehicle charging, HVAC control, and long-term planning for renewable energy installations. This novel strategy represents a significant advancement in context-based optimization using LLMs, facilitating sustainable energy practices customized to individual needs.

I. INTRODUCTION

Despite computational advances, the complex challenges of meeting energy demands and reducing carbon emissions persist. Optimization techniques, as seen in EV charging [1], energy storage [2], renewable investments [3], smart building operations [4], and demand side management [5], hold great promise. **However, the broader democratization of these tools remains a significant hurdle.** This paper tackles this issue, advocating for the accessibility and practicality of advanced computational methods for all, especially the underserved [6]. The emergence of LLMs offers a breakthrough in overcoming these barriers. We demonstrate that LLMs can bridge the gap between the high costs of traditional optimization and the need for personalized, accessible solutions. Our goal is not industry-grade optimization, but rather to provide users with tools for informed decision-making, potentially involving optimization problem formulation. This approach leverages LLMs to streamline modeling and optimization, enabling interaction through natural language without extensive programming or mathematical optimization knowledge.

Technical challenges and solutions. While LLMs such as OpenAI’s ChatGPT offer a natural substrate for conversational AI (a technology that enables machines to understand, interpret, and engage in natural conversations), it is essential to recognize that current LLMs have not excelled in tasks like arithmetic and logical reasoning [7]. Recent works have introduced methods such as chain of thoughts [8] (and variants such as algorithm of thoughts [9]), which prompts a series of intermediate reasoning steps, and autoformalism [10], which automatically translates natural language mathematics to formal specifications and proofs. However, many energy problems, like energy storage control and long-term planning for PV panel installation, demand complex decision-making. These problems differ from arithmetic reasoning, common-sense reasoning, and symbolic reasoning in the following aspects: *Complexity*, as they often involve numerous vari-




Fig. 1: Comparison of the baseline method (simple prompting with GPT-4) and EC using either CVXPY or SCIPY as the code projection layer across various problem categories. Success rate denotes the percentage of instances in which the provided answer aligns with user expectations, satisfies all given constraints, and is either close to or exactly optimal.

ables, constraints, and objectives with potentially nonlinear relationships between variables; and *incomplete information*, as energy systems are influenced by various factors, such as user preferences, which may not be provided in the initial problem description.

A. Contributions

In this research, we address the above challenges as follows. To manage complexity, we develop a procedure to map a problem description to code using the grammar of packages that support optimization formulation, and to iteratively formulate, debug, and execute the program. We also introduce external tooling to execute the written code and solve the formulated optimization using dedicated algorithms. Furthermore, we implement prompt engineering techniques to enable LLMs to accurately understand and respond to user-specific preferences. To address incomplete information, we leverage the reasoning capabilities of LLMs to identify key parameters and use a question-and-answer format in natural conversations to solicit this information from users. Additionally, we employ LLMs for auto-informalism to explain the solution to users. An overview of the EC framework is depicted in Fig. 2. This approach bears dual benefits: it not only guides the user through complex nonlinear reasoning tasks of energy saving planning but also offers an insightful explanation of the solution.

We further demonstrate the possibility of our proposed approach in solving a variety of tasks within the energy domain, ranging from EV charging, HVAC and battery control, to long-term planning problems such as cost-benefit evaluation of installing rooftop solar PVs, heat pumps, and battery sizing. In comparison to the simple prompting method, where the task description is directly presented to the LLM,

our proposed EC framework improves the success rate, as illustrated in Fig. 1. Although the optimizations formulated are not exceedingly sophisticated, we find that the LLM effectively addresses the problem to a considerable degree. Throughout the remainder of the paper, unless specified otherwise, we refer to OpenAI’s ChatGPT (GPT-4) when mentioning LLM.

B. Related Work

Recent work has explored using LLMs for mathematical autoformalization, converting natural language into formal languages like Isabelle/HOL [11]. Analogously, we use intermediaries like SCIPY and CVXPY to transform intuitive formulations of optimization problems into standard form [12]. Despite few CVXPY examples online, LLMs exhibit surprising adeptness, albeit with occasional syntax errors that we correct via Python debugging. While LLMs struggle with multi-step problems [13], techniques like in-context learning [14] and algorithmic search [9] have shown promise. We present a novel viewpoint using optimization itself as a reasoning tool to enrich existing LLM augmentation techniques [15]. LLMs can also revise responses given feedback, e.g. in question-answering and code debugging [16]. We incorporate revisions based on programming language feedback and user requests, enabling rapid error correction.


Optimizing energy systems requires commonsense reasoning and domain knowledge [17]. Prior work has used ML and expert systems [18], but we harness LLM expertise for optimization through natural language interaction. To our knowledge, we are the first to tackle non-trivial, energy-specific problems with incomplete information this way.

The remaining sections of this paper are structured as follows. Section II describes our proposed framework in detail, including the design principles and the optimization autoformalism approach. Section III presents experimental results that demonstrate the effectiveness of our approach in solving various energy-related tasks. We discuss the potential and future directions of conversational AI for energy sustainability and conclude in Section IV.

II. ENERGY CONCIERGE FRAMEWORK

Using an EV charging query as an example, our Energy Concierge operates as follows:

- 1) User submits natural language request (e.g. optimize my charging schedule). LLM determines if it is an optimization problem.
- 2) If so, LLM requests input parameters needed to solve it. It may ask clarifying questions (e.g. charging capacity? preferred hours?).
- 3) LLM formulates Python code with user information, translating the query into an optimization problem via autoformalism.
- 4) Interface executes code to solve problem, with debugging iterations if needed.
- 5) LLM explains the optimal solution clearly (charging schedule, cost savings). Enabling informed user decisions.



Step annotations

1a, 1b: problem statement, prompt engineered input	4a, 4b: LLM provide python code, program execution
2a, 2b: LLM’s response, asking user input parameters	5a, 5b: Code output, asking LLM to explain the result
3a, 3b: User details, passing values to LLM	6a, 6b: (Optional) debugging
	7a, 7b: LLM explanation

Fig. 2: Energy Concierge framework. The user engages with an LLM through natural language queries and responses. The LLM identifies the necessary input parameters for optimization and generates Python code to address the problem. The program interface then executes the code and relays the solution back to the LLM, which subsequently provides a clear explanation to the user.

Through interactivity and personalization grounded in optimization and autoformalism, this conversational framework empowers users to improve energy efficiency.

A. Optimization autoformalism

A general optimization problem can be written as:

$$\begin{aligned}
 &\text{minimize}_{x \in \mathbb{R}^p} && f(x; \theta) \\
 &\text{subject to} && g_i(x; \theta) \leq 0, \quad i = 1, 2, \dots, m \quad (P(\theta)) \\
 &&& h_j(x; \theta) = 0, \quad j = 1, 2, \dots, n,
 \end{aligned}$$

where we seek to minimize the objective function $f(x; \theta)$ subject to a set of inequality constraints $g_i(x; \theta) \leq 0$, and equality constraints $h_j(x; \theta) = 0$. Here, $x \in \mathbb{R}^p$ is the decision variable and θ is the collection of hyperparameters that defines the optimization instance, including objective and constraint functions; in other words, the solution of the optimization $P(\theta)$ can be regarded as a function of the hyperparameters θ [19].

Optimization techniques can be used to solve energy-related problems [1]–[5]. Automating the formulation and solution of optimization problems is essential due to the technical skills gap, challenges in manually incorporating user preferences and constraints, and the inefficiencies in manual modifications based on user feedback. The human-guided autoformalism proposed in this study automatically translates natural language task specifications to optimization instances. However, directly implementing this approach may face issues such as ambiguity, incompleteness, and incorporating user-specific preferences. The subsequent subsections will delineate strategies to address them.

1) Optimization formulation: The optimization process begins with identifying the objective function, decision variables, and constraints from the user’s task description [T]. We tested two approaches:

- Approach 1: Directly ask the LLM to identify the optimization components in [T].
- Approach 2: First prompt the LLM to identify the 5 most important parameters in [T], then use these parameters to formulate the optimization instance.

Approach 2 outperforms Approach 1 because LLMs like GPT-4 are trained on diverse data, not just optimization tasks. Asking for the top 5 parameters simplifies the task and provides context, helping LLMs generate more precise optimization formulations [8], [9]. Directly requesting optimization formulations (Approach 1) can introduce ambiguity; initial parameter identification (Approach 2) reduces this uncertainty.

After identifying the essential parameters, the natural language query is transformed into a computational instance using the prompt: “Write a Python code using [lib] to solve this optimization problem,” where [lib] is either CVXPY or SCIPY. We found that SCIPY yields a higher success rate (71% to 100%) compared to CVXPY (51% to 80%) (see Fig. 1). This difference is likely due to SCIPY’s broader acceptance and longer existence in the Python community, resulting in more SCIPY-related content available online for the model to leverage during training.¹

2) *Solving an optimization and debugging*: To address LLMs’ limitations in nonlinear reasoning [7], we enable the model to interact with an external Python program to solve optimization tasks. We extract the code block from the LLM’s output using regular expressions, searching for unique delimiters enclosing the code. This method reliably extracts code without requiring a large labeled dataset, unlike training a machine learning model [20]. During development, we encountered two types of errors in the LLM-generated code: erroneous translation into an optimization problem and syntactic bugs in an otherwise correctly translated problem.²

To rectify translation errors, we rely on user interaction and clarification of the formulated optimization. For syntactic errors, we use an automated process: identify the error message, feed it to the LLM to isolate relevant code snippets, generate potential remedies, and assess the proposed solutions in a new iteration (Fig. 7 in Appendix). Taking multiple code samples before debugging often results in fewer LLM queries (Fig. 5). By following this procedure, we successfully resolved most errors encountered during our experiments.

3) *Optimization auto-informalism*: After discovering the optimal solution, the LLM articulates the results in a comprehensible, natural language format. This system provides detailed explanations of the optimal solution and any constraints or preferences factored in during the optimization process (Steps 7a and 7b in Fig. 2). This auto-informalism approach complements autoformalism by offering intuitive

¹A search on GitHub (as of 4/25/24) returns 8K repositories for “SCIPY” and only 272 for “CVXPY,” supporting our preference for SCIPY.

²Errors can also arise due to the limitations of the optimization solvers, but since most energy problems involve linear or quadratic convex optimization, we presume such errors are comparatively infrequent. We do not classify infeasibility as an error.




Fig. 3: A comparison of EV charging plans to improve the cost and load on the grid. Observations reveal that EC effectively capitalizes on hours with lower prices in contrast to the baseline.

insight into the optimization outcomes.

III. EXPERIMENTS AND CASE STUDIES

This section examines the potential of LLM-based autoformalism in real-time decision-making (Sec. III-A) and sustainable long-term planning (Sec. III-B). A general analysis of the proposed methods is provided in Sec. III-C. The [21, appendix] offers example interactions.

A. Real-time decision making

1) *Smart EV charging*: The core issue in smart EV charging is to optimize EV charging patterns to balance power grid loads, mitigate energy costs, and fulfill user preferences [1]. EC can correspond user-provided information to optimization parameters, such as the maximum charging rate. It grasps the user’s charging availability and translates this information into decision variables and constraints, leading to an optimized schedule (see Fig. 3).

2) *HVAC control*: HVAC control is a pivotal issue that has been extensively studied [4]. However, the algorithms that have been developed are often more advanced than the current control panels in most buildings, leading to a disconnect between sophisticated methods and user comprehension. The heart of the issue is to discover a setpoint that ensures a comfortable indoor climate while minimizing energy consumption. EC provides clear and actionable advice like “Set your thermostat to the optimal temperature (75°F in this case) during hot, humid days...” with a clear rationale. It also offers customized suggestions based on user conditions, along with valuable energy efficiency tips such as “monitor your energy consumption”, and “adapt to changing conditions”. Nevertheless, the cost of simplicity is the inability to perform sophisticated controls like pre-heating/cooling, which entails heating the room in anticipation of occupancy and can only be accounted for through a multiperiod formulation involving intricate room thermal dynamics [4].

3) *Battery charging control*: This problem involves optimizing the charging and discharging cycles of a home battery system, considering factors such as electricity pricing, solar generation, and household demand, as seen in competitions like the CityLearn Challenge [22]. As discussed in [23], selecting the appropriate optimization problem for a given context is crucial. Unlike [23], where the context is derived

from a reward signal, our context is provided by the user in natural language. The EC interprets this context accurately to establish the right optimization problem.

EC successfully formulates an accurate objective that computes the weighted sum of total electricity cost, considering the electricity price. It also correctly formulates the lower and upper limits for the charging and discharging variables, the temporal interdependence of the state of charge and the charging rate, and the maximum capacity constraint. The energy balance constraint identified demonstrates LLM’s understanding of world modeling, including this constraint *without* specific user instruction. EC’s explanation further shows an understanding of physical constraints and its overall objective. However, EC implicitly assumes that excess energy produced each hour can be sold back to the grid at the same rate as it was purchased, which is not typically the case. Future work can explore methods to encourage EC to be more explicit about assumptions when constructing optimization instances.

B. Long-term planning for sustainability

1) Cost-benefit analysis of installing rooftop solar PVs:

To perform a cost-benefit analysis of solar PV, one must estimate the costs and benefits over the system’s lifespan and compare them to a relevant alternative system [24].

EC shows proficiency in understanding physical constraints, such as the stipulation that the panel area shouldn’t surpass the roof area, and proposes that the installed area be adequate to supply the electricity demand. However, it operates under a few assumptions, such as the user planning to source all electricity demand from PV to fully leverage the budget, and that the PV’s efficiency stands at 0.12. While these assumptions are generally sensible, stating them in the explanations would make the model’s workings more transparent. Additionally, the model seems to underutilize the user’s provided information, such as the building’s location, which could potentially inform the required area-to-power conversion efficiency more accurately. This is understandable given the model’s lack of access to an external database, and future work that enhances this capability could be a worthwhile pursuit.

2) Cost-benefit analysis of installing a heat pump:

Key considerations in the cost-benefit analysis of a heat pump include the initial cost of installation, which can vary based on the type, size, and complexity of the system, as well as the availability of incentives and rebates [25].

Similar to the PV installation scenario in Sec. III-B.1, EC proficiently identifies pivotal relationships, including the annual operating cost of a central AC and a heat pump. EC’s approach focuses on optimizing annual savings, leaving the actual estimation of the payback period to the user. For instance, EC calculates that the annual savings with a heat pump equals \$550, then elaborates on how to use this data for investment decisions, suggesting “If the purchase and installation of the heat pump cost \$5,000, the payback period would equate to around 9.1 years ($\$5,000 / \550).”




Fig. 4: The efficiency of EC when integrated with SCIPY is gauged by the compilation success of the Python code it generates, the solution’s precision, and the lucidity of the explanation presented to the user. The black lines highlight the performance shifts observed when employing CVXPY.

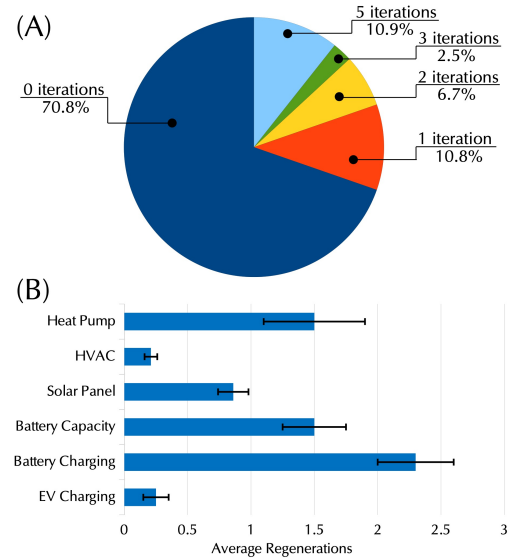


Fig. 5: (A) A comparative analysis of the number of iterations needed to achieve executable code across all problem categories (note that none of the experimental instances use 4 iterations). (B) A comparison of the average number of code generations needed for EC to produce executable code, limited to a maximum of 5. The black bars represent the variation in iterations across problem categories over 20 random runs.

EC also suggests users compare alternatives, plan long-term strategies, evaluate environmental impact, and seek incentives or rebates. This shows that EC can provide not only a decision but also the context and information necessary to empower users to make more cost-effective and environmentally conscious choices.

C. General findings and analysis

To elucidate the types of errors EC commits, we display the success rates concerning error-free code production, logical correctness of the code, and the ability to articulate the solution derived from the optimization code to the end user in Fig. 4. Our framework enables EC to achieve a 100% compilation rate. When EC generates an error-free optimiza-

tion code, it consistently explains the results to the end user. SCIPY outperforms CVXPY in overall performance, aligning with our observation that SCIPY has broader online recognition and use, where most LLMs receive training.

Fig. 5 presents the average number of code regenerations needed and the average debugging iterations required to obtain syntax error-free optimization code. Although each subproblem within the optimization problem category exhibits different success rates, a relatively low number of code regenerations are typically needed. When the LLM consistently commits errors, incorporating debugging information in our framework is essential for achieving optimization solutions. Reducing the number of generated code samples from $s = 5$ to $s = 1$ slightly increases the required debugging iterations across all examples, suggesting that debugging is a more efficient error-handling strategy compared to simple regeneration.

Estimating the Probability of One-Round Autoformalism Success. The aforementioned results allow us to infer the probability of generating valid code successfully, denoted as p . We postulate that the event of successfully generating valid code is independent and identically distributed for the same category of problems, following a truncated geometric distribution with a cap at 5 trials. Based on this supposition, we deduce that $\sum_{k=1}^5 k(1-p)^{k-1}p + 6 \sum_{k \geq 6} (1-p)^{k-1}p = z$, where z signifies the number of generations needed to achieve success (with generations capped at 5, as shown in Fig. 5), and p is the probability of success to be estimated. Substituting the values for z from Fig. 5 into the equation and solving for p , we obtain the one-round autoformalism success rates $p = 0.8, 0.25, 0.38, 0.53, 0.83, 0.38$ for the problems of EV Charging, Battery Charging, Battery Capacity, Solar Panel, HVAC Control, and Heat Pump Investment, respectively.

Estimation of the Probability of Debugging Success. We posit that the event of successful debugging, given the presence of an erroneous code, is independent and identically distributed across all problem classes. Using Fig. 5, we first normalize the frequency of the required number of debugging iterations by the frequency of generating an erroneous code in the first run (i.e., $1 - 0.7 = 0.3$). This reveals that the frequencies of observing debugging iterations 1, 2, 3, 4, and ≥ 5 are represented by y_k , where $k \in \{1, 2, \dots, 5\}$. Let q denote the probability of successfully debugging the code. From this, we can construct a system of polynomial equations in the form of $(1-q)^{k-1}q = y_k$ for $k \in \{1, 2, 3, 4\}$, and $1 - \sum_{k=1}^4 (1-q)^{k-1}q = y_5$ given that debugging is capped at 5 iterations. Using a line search between 0 and 1 in increments of 0.01, we calculate the values of \hat{y}_1 to \hat{y}_5 for each q using the specified equations, and determine the mean squared error (MSE) between the vectors \hat{y} and y . Our optimal estimate for the probability is $q = 0.26$.

Optimality Gap and Improvement Over Baseline. The optimality gap for a feasible solution is calculated as the ratio $v/v^* - 1$, where v denotes the objective value of the candidate solution and v^* signifies the optimal value of the corresponding minimization problem. In Fig. 4, we classify test




Fig. 6: We assess the average optimality gap by comparing the exact solution of the optimization problem with the solution determined by EC, alongside evaluating the enhancement over the baseline in terms of optimality.

cases as correct only if they equate to the globally optimal solution value, i.e., a 0% optimality gap. Additional results presented in Fig. 6 demonstrate the average optimality gap when incorrect logic within the code results in the omission of key parameters (such as the efficiency of a battery). Our framework was evaluated using 20 examples for each energy optimization problem, including EV charging. As depicted, the instances display an optimality gap generally within the 20% range, with the majority under 10%.

For the baseline method, the model was simply prompted with our question along with necessary parameters to solve the optimization problem. We noticed that even when the problem is clearly an optimization issue, LLMs may opt to generate responses by attempting to apply logic towards reaching an answer. As seen in Fig. 6, this approach does not yield favorable results in comparison to the EC framework. This can be attributed to the fact that many energy-related optimization problems of importance do not yield closed-form solutions, making the correct utilization of convex program solvers critical.

We further report on the improvement over the baseline, measured by $v_b/v - 1$, where v_b represents the objective value of the baseline solution. As is evident, the improvement can amount to as much as 60%, with most instances falling within the 30% range. Some of the problems yielding the greatest improvements include Solar Panel, Heat Pump, Battery Charging, and EV Charging, which encompass both real-time decisions and long-term investments.

IV. DISCUSSIONS AND CONCLUSION

Conversational AI for sustainability has long seemed aspirational, but LLMs offer new potential. By providing an intuitive interface, our framework helps individuals engage more consciously about energy use. Greater awareness drives sustainable behaviors like responsible consumption, efficiency, and solar adoption [26], [27].

Households average 3.5 kWh peak electricity use, while EV chargers approach 10 kWh. EV charging at peak triples load, straining infrastructure. Utilities offer off-peak rates, but financial incentives alone don't fully shift behaviors—an "efficiency gap" phenomenon [26], [27]. Our framework simplifies this via conversation. With 2022's 3 million EVs, adopting off-peak charging could save \$876 million annually (at \$0.06 vs \$0.14 per kWh rates) [28], [29]. With 2035's

projected 73 million EVs [30], savings could reach \$21.3 billion. Beyond consumer savings, off-peak charging lessens grid stress, stabilizes prices, and reduces needs for new plants. By simplifying optimization, our system enables impactful sustainability actions.

A key limitation of EC is its reliance on the LLM's ability to accurately formulate optimization problems, ensuring proper alignment of objectives and constraints. Incorrect formulations can lead to ineffective or harmful solutions, despite the perceived reliability of these systems. Auto-informalism (Sec. II-A.3) helps mitigate this risk by scrutinizing problem formulations more carefully and alerting users to potential issues or assumptions. Enhancing auto-informalism's effectiveness is thus critical for reliable solutions. Additional reliability methods like validation routines, robustness testing, and expanded user feedback loops can further strengthen the framework.

While we merely sketch out the potential, the proposed shift towards human-guided optimization autoformalism could democratize access to sophisticated technologies and set the stage for a more equitable and sustainable future.

REFERENCES

- [1] S. M. Arif, T. T. Lie, B. C. Seet, S. Ayyadi, and K. Jensen, "Review of electric vehicle technologies, charging methods, standards and optimization techniques," *Electronics*, vol. 10, no. 16, p. 1910, 2021.
- [2] R. Khezri, A. Mahmoudi, and H. Aki, "Optimal planning of solar photovoltaic and battery storage systems for grid-connected residential sector: Review, challenges and new perspectives," *Renewable and Sustainable Energy Reviews*, vol. 153, p. 111763, 2022.
- [3] Y. Yang, S. Bremner, C. Menictas, and M. Kay, "Battery energy storage system size determination in renewable energy systems: A review," *Renewable and Sustainable Energy Reviews*, vol. 91, pp. 109–125, 2018.
- [4] J. Drgoňa, J. Arroyo, I. C. Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie *et al.*, "All you need to know about model predictive control for buildings," *Annual Reviews in Control*, vol. 50, pp. 190–232, 2020.
- [5] B. P. Esther and K. S. Kumar, "A survey on residential demand side management architecture, approaches, optimization models and methods," *Renewable and Sustainable Energy Reviews*, vol. 59, pp. 342–351, 2016.
- [6] J. Currie, D. I. Wilson, N. Sahinidis, and J. Pinto, "Opti: Lowering the barrier between open source optimizers and the industrial matlab user," *Foundations of computer-aided process operations*, vol. 24, p. 32, 2012.
- [7] J. W. Rae, S. Borgeaud, T. Cai, K. Millican, J. Hoffmann, F. Song, J. Aslanides, S. Henderson, R. Ring, S. Young *et al.*, "Scaling language models: Methods, analysis & insights from training gopher," *arXiv preprint arXiv:2112.11446*, 2021.
- [8] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. H. Chi, Q. V. Le, D. Zhou *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," in *Advances in Neural Information Processing Systems*, 2022.
- [9] B. Sel, A. Al-Tawaha, v Vanshaj, R. Jia, and M. Jin, "Algorithm of thoughts: Enhancing exploration of ideas in large language models," *ICML*, 2024.
- [10] Y. Wu, A. Q. Jiang, W. Li, M. N. Rabe, C. E. Staats, M. Jamnik, and C. Szegedy, "Autoformalization with large language models," in *Advances in Neural Information Processing Systems*, 2022.
- [11] Y. Wu, A. Q. Jiang, W. Li, M. Rabe, C. Staats, M. Jamnik, and C. Szegedy, "Autoformalization with large language models," *Advances in Neural Information Processing Systems*, vol. 35, pp. 32 353–32 368, 2022.
- [12] A. Agrawal, R. Verschueren, S. Diamond, and S. Boyd, "A rewriting system for convex optimization problems," *Journal of Control and Decision*, vol. 5, no. 1, pp. 42–60, 2018.
- [13] J. Austin, A. Odena, M. Nye, M. Bosma, H. Michalewski, D. Dohan, E. Jiang, C. Cai, M. Terry, Q. Le *et al.*, "Program synthesis with large language models," *arXiv preprint arXiv:2108.07732*, 2021.
- [14] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, "Chain-of-thought prompting elicits reasoning in large language models," *Advances in Neural Information Processing Systems*, vol. 35, pp. 24 824–24 837, 2022.
- [15] G. Mialon, R. Dessì, M. Lomeli, C. Nalmpantis, R. Pasunuru, R. Raileanu, B. Rozière, T. Schick, J. Dwivedi-Yu, A. Celikyilmaz *et al.*, "Augmented language models: a survey," *arXiv preprint arXiv:2302.07842*, 2023.
- [16] T. X. Olausson, J. P. Inala, C. Wang, J. Gao, and A. Solar-Lezama, "Demystifying gpt self-repair for code generation," *arXiv preprint arXiv:2306.09896*, 2023.
- [17] A. B. Saka, L. O. Oyedele, L. A. Akanbi, S. A. Ganiyu, D. W. Chan, and S. A. Bello, "Conversational artificial intelligence in the aec industry: A review of present status, challenges and opportunities," *Advanced Engineering Informatics*, vol. 55, p. 101869, 2023.
- [18] R. Panchalingam and K. C. Chan, "A state-of-the-art review on artificial intelligence for smart buildings," *Intelligent Buildings International*, vol. 13, no. 4, pp. 203–226, 2021.
- [19] M. Jin, V. Khattar, H. Kaushik, B. Sel, and R. Jia, "On solution functions of optimization: Universal approximation and covering number bounds," *AAAI*, 2023.
- [20] S. Iyer, I. Konstas, A. Cheung, and L. Zettlemoyer, "Summarizing source code using a neural attention model," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2016, pp. 2073–2083.
- [21] M. Jin, B. Sel, H. FNU, and W. Yin, "Democratizing energy management with llm-assisted optimization autoformalism," 2024, full Version: <http://www.jinming.tech/papers/ChatEnergy23-smartgridcomm.pdf>.
- [22] J. R. Vázquez-Canteli, J. Kämpf, G. Henze, and Z. Nagy, "Citylearn v1.0: An openai gym environment for demand response with deep reinforcement learning," in *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, ser. BuildSys '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 356–357. [Online]. Available: <https://doi.org/10.1145/3360322.3360998>
- [23] V. Khattar and M. Jin, "Winning the citylearn challenge: Adaptive optimization with evolutionary search under trajectory-based guidance," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 37, no. 12, 2023, pp. 14 286–14 294.
- [24] M. Thebault and L. Gaillard, "Optimization of the integration of photovoltaic systems on buildings for self-consumption—case study in france," *City and Environment Interactions*, vol. 10, p. 100057, 2021.
- [25] F. Belaïd, Z. Ranjbar, and C. Massié, "Exploring the cost-effectiveness of energy efficiency implementation measures in the residential sector," *Energy Policy*, vol. 150, p. 112122, 2021.
- [26] O. I. Asensio and M. A. Delmas, "Nonprice incentives and energy conservation," *Proceedings of the National Academy of Sciences*, vol. 112, no. 6, pp. E510–E515, 2015.
- [27] T. D. Gerarden, R. G. Newell, and R. N. Stavins, "Assessing the energy-efficiency gap," *Journal of economic literature*, vol. 55, no. 4, pp. 1486–1525, 2017.
- [28] APPALACHIAN POWER, "Virginia s.c.c. tariff no. 26 appalachian power company," 2021. [Online]. Available: <https://www.appalachianpower.com/lib/docs/ratesandtariffs/Virginia/Tariff26-MASTER-Standard-June1-2023RPS-RAC.pdf>
- [29] IEA, "Trends in electric light-duty vehicles," 2023. [Online]. Available: <https://www.iea.org/reports/global-ev-outlook-2023/trends-in-electric-light-duty-vehicles>
- [30] NREL, "National economic value assessment of plug-in electric vehicles," 2016. [Online]. Available: <https://www.nrel.gov/docs/fy17osti/66980.pdf>
- [31] H. C. Hesse, R. Martins, P. Musilek, M. Naumann, C. N. Truong, and A. Jossen, "Economic optimization of component sizing for residential battery storage systems," *Energies*, vol. 10, no. 7, p. 835, 2017.

APPENDIX

A. Debugging Illustration

Fig. 7 illustrates the process of solving the optimization problem and debugging the LLM-generated code. The flow follows the **No Errors** path if the code executes seamlessly; otherwise, it diverts to the **Execution Errors** path. Debugging information is incorporated into a refined prompt for the LLM to generate a new code iteration. This iterative process continues until the code runs without errors, ensuring the optimization problem is correctly translated and syntactically sound.

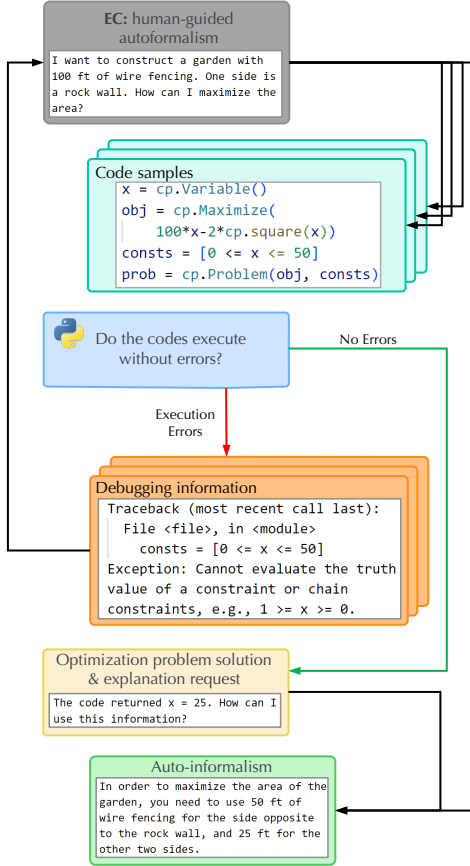


Fig. 7: Illustration of the process of solving the optimization problem and debugging the LLM-generated code.

B. Additional examples and interaction history

1) HVAC control:

User: Is it possible to save money by adjusting the HVAC temperature?

EC: Yes, adjusting the HVAC temperature can help save money on energy costs. To provide you with more specific recommendations, please provide the following information:

- Indoor temperature setting: What is your preferred temperature setting for both heating and cooling? Knowing your current setpoint will help determine if adjustments can be made for energy savings.
- Occupancy schedule: What is your typical daily schedule for when people are at home and away? This information will help identify potential opportunities for adjusting the temperature when the building is unoccupied.
- Type of HVAC system: What type of heating and cooling system do

you have (e.g., central air conditioning, heat pump, furnace, etc.)? The efficiency and potential energy savings may vary depending on the system type.

- Insulation and building envelope: How well insulated is your home or building? Please provide information about the insulation in the walls, attic, and windows. This will help identify any areas where improvements can be made to further reduce energy consumption.
- Local climate: What is the typical climate in your region (e.g., mild, cold, hot, humid, etc.)? This information will help tailor recommendations based on the specific climate conditions where you live.

User: (note that the user directly answers the questions)

- preferred temperature is 65-75
- full day occupancy
- heat pump
- average insulation
- hot and humid climate

Background processing:

$$\text{minimize}_{C \in \mathbb{R}} C \quad (1a)$$

$$\text{subject to } C = c \frac{|x - y|}{e} z, \quad (1b)$$

$$h \leq x \leq k, \quad (1c)$$

where C , c , x , y , e , z , h , k represent total cost, electricity cost per kWh, adjusted temperature, ambient temperature, heat pump efficiency, occupancy, heating setpoint and cooling setpoint.

```

# Code (only given with important lines)
adj_temp = cp.Variable() # HVAC adjusted indoor temperature
constraints = [
    adj_temp >= heating_setpoint,
    adj_temp <= cooling_setpoint,
]
temp_difference = cp.abs(adj_temp - ambient.temp)
energy_cost = elec_cost_per_kwh * (temp_difference / heat_pump_efficiency) * occupancy
objective = cp.Minimize(energy_cost)
problem = cp.Problem(objective, constraints)
problem.solve()
  
```

Output:
The optimal adjusted temperature: 75.00
The minimum energy cost per day: 9.60

EC: The optimal adjusted temperature of 75°F provided by the CVXPY optimization indicates that, under the given parameters and constraints, setting your indoor temperature to 75°F would minimize your energy cost for running the heat pump.

Here's how you can use this information:

- Adjust your thermostat: Set your thermostat to the optimal temperature (75°F in this case) during hot, humid days. This will help you save energy by reducing the amount of cooling your heat pump needs to provide.
- Programmable thermostat: If you have a programmable thermostat, you can schedule temperature adjustments to coincide with your daily occupancy pattern. For example, you can set the thermostat to 75°F during peak usage hours and adjust it to a higher temperature when the building is unoccupied. This can help you save even more on energy costs.
- Monitor your energy consumption: Keep track of your energy consumption over time to verify the savings achieved by implementing these changes. Compare your energy bills before and after the adjustments to see if there's a noticeable reduction in costs.
- Adapt to changing conditions: Keep in mind that the optimal temperature may change depending on the outdoor temperature, humidity, and other factors. You can periodically re-run the optimization with updated parameters to find the most energy-efficient temperature settings for different conditions throughout the year.
- Additional energy-saving measures: While adjusting the temperature is a good first step, there are other ways to improve your home's energy efficiency. You can consider upgrading your insulation, sealing gaps and cracks in the building envelope, installing energy-efficient windows, and using shading devices like awnings or blinds to reduce solar heat gain.
- Remember that the optimal temperature calculated in this example is based on the specific parameters provided. To get more accurate results tailored to your situation, you may need to adjust the input values, such as electricity costs, heat pump efficiency, or local climate conditions.

2) Battery Charging Control:

User: I live in California. I have a battery of 20kwh. And solar panels on the roof with max 3kw electricity production. My electricity price is \$0.2 between 6pm and 10pm, \$0.1 otherwise. Assume some random electricity production and consumption in the house for the next 24 hours. Can you write me a CVXPY code to optimize for cost by controlling the battery charging or discharging? Max charge or discharge rate of the battery is 3kw.

Background processing:

$$\begin{aligned} \text{minimize}_{x \in \mathbb{R}^{24}} \quad & \sum_{t=1}^{24} c_t P_{e,t} \quad (2a) \\ \text{subject to} \quad & -P \leq x_t \leq P, \quad t = 1, 2, \dots, 24 \quad (2b) \\ & b_0 = 0, \quad (2c) \\ & b_t = b_{t-1} + x_t, \quad t = 1, 2, \dots, 24 \quad (2d) \\ & 0 \leq b_t \leq B, \quad t = 1, 2, \dots, 24 \quad (2e) \\ & c_t = x_t - s_t + d_t, \quad t = 1, 2, \dots, 24, \quad (2f) \end{aligned}$$

where c, b, x, P, s, d represent the electricity usage, battery state, battery charging rate, maximum power of the battery, solar energy production and demand of the household, respectively.

```
# Code (only given with important lines)
battery_capacity = 20 # kWh
max_power = 3 # kW
price_peak = 0.2 # $/kWh
price_off_peak = 0.1 # $/kWh

np.random.seed(0)
solar_production = np.random.uniform(0, 0.5, 24)
household_consumption = np.random.uniform(0,
max_power, 24)
battery_power = cp.Variable(24)
battery_state = cp.Variable(24)
cost = cp.Variable(24)
objective = cp.Minimize(sum(cost))
constraints = [battery_power <= max_power,
battery_power >= -max_power,
battery_state >= 0,
battery_state <= battery_capacity,
# Battery state evolves over time
battery_state[0] == 0, # Assume the battery
starts empty
battery_state[1:] == battery_state[:-1] +
battery_power[1:],
battery_state[-1] == 0,
]
for i in range(24):
constraints.append(cost[i] == (battery_power[i] -
solar_production[i] + household_consumption[i])
* (price_peak if 18 <= i < 22 else
price_off_peak))
prob = cp.Problem(objective, constraints)
prob.solve()
# Output:
Total cost: 1.6952066412652915
Battery power (kW): [-3.00000056 2.2385663
1.72916783 2.15337281 2.39723956 0.91568419
0.34981417 0.13375868 0.05146217 0.02062781
0.0104214 0.01063623
0.0214874 0.053826 0.13999059 0.36614605
0.95844804 2.50919925
-2.9999989 -2.99997493 -3.0000141 -3.00001957
-1.01609253 -1.04374843]
```

EC: The optimization program provided generates two main outputs:

- Total cost: This is the total estimated cost of your electricity consumption for the given 24 hour period, under the optimal charging and discharging policy of the battery. In this case, your total cost is estimated to be approximately \$1.70.
- This array shows the optimal charge (-) or discharge (+) rate of the battery for each hour in the 24-hour period. For example, in the first hour (-3.00000056 kW), the battery should be charged at nearly the maximum rate. In the second hour (2.2385663 kW), the battery should be discharged at roughly 2.24 kW. This indicates that the solar panels are likely generating less power than the household is consuming. This charging and discharging strategy continues throughout the 24-hour period to minimize the cost of electricity.

3) Cost-benefit analysis of installing rooftop solar PVs:

User: Should I install solar panels on my roof? How much can it save me?

EC: Installing solar panels on your roof can be a great way to save on electricity costs and reduce your carbon footprint. However, the amount you can save depends on various factors. To help me provide an accurate estimate of the potential savings, please provide the following information:

- Location: In which city or region is your home located? This will help determine the amount of sunlight your solar panels can receive.
- Roof size and orientation: What is the size of your roof and its orientation (e.g., facing south, east, or west)? This will affect the number of solar panels that can be installed and their efficiency.
- Energy consumption: What is your average monthly electricity consumption in kilowatt-hours (kWh)? This will help determine the size of the solar panel system required to meet your energy needs.
- Local electricity rates: What is the cost per kWh for electricity in your area? This will help calculate your potential savings from using solar power.
- Budget: What is your budget for installing a solar panel system? This will help determine the size and quality of the system you can install, which will affect your potential savings.

User: (note that the user directly answers the questions)

- Seattle USA
- 300 sq ft facing east
- 400 kWh per month
- USD 0.13 per kWh
- 8000

Background processing:

$$\text{minimize}_{U \in \mathbb{R}} \quad U \quad (3a)$$

$$\text{subject to} \quad U = B - p \times A, \quad (3b)$$

$$R \geq A, \quad (3c)$$

$$A \times w \times 0.12 \geq M, \quad (3d)$$

$$A \geq 0, \quad (3e)$$

where U, B, p, A, R, w, M represent unused budget, budget, panel price rate, panel area, roof area, panel wattage per sqft and monthly consumption.

```
# Code (only given with important lines)
panel_area = cp.Variable()
objective = cp.Minimize((budget - (panel_area * 10)))
constraints = [
panel_area <= roof_area, # Panel area should not
exceed the roof area
panel_area * panel_wattage_per_sqft * 0.12 >=
monthly_consumption, # Monthly energy
production should meet or exceed consumption
panel_area >= 0 # Panel area should be non-
negative
]
problem = cp.Problem(objective, constraints)
result = problem.solve()
annual_savings = 12 * monthly_consumption *
electricity_rate - (12 * (panel_area.value *
panel_wattage_per_sqft * 0.12 -
monthly_consumption) * electricity_rate)
# Output:
Optimal panel area (sq ft): 300.000000102021
Total cost (USD): 3000.00000010202098
Monthly energy production (kWh): 540.0000001836378
Annual savings (USD): 405.59999971352505
```

EC: This information can be useful for someone considering installing solar panels on their roof. They can use the results to make an informed decision about whether or not to invest in solar panels and what size of installation would be optimal for their situation. The results provide the following insights:

- Optimal panel area (sq ft): This is the recommended size of the solar panel installation that satisfies the given constraints, such as roof area and monthly energy consumption. In this case, it is nearly the maximum possible area (300 sq ft).
- Total cost (USD): This is the estimated cost of the solar panel installation based on the optimal panel area. In this case, the cost is \$2,999.99, which is within the budget of \$8,000.
- Monthly energy production (kWh): This is the estimated amount of electricity the solar panels will produce each month, based on the

optimal panel area. In this case, it is 539.99 kWh, which is more than the monthly consumption of 400 kWh.

- Annual savings (USD): This is the estimated amount of money the solar panels will save on electricity bills over the course of a year. In this case, the annual savings are \$405.60.

Using this information, someone can determine if investing in solar panels makes financial sense for their situation. They can weigh the cost of the installation against the expected savings and consider other factors, such as environmental benefits or local incentives for renewable energy. Additionally, they can use the optimal panel area as a starting point for discussions with solar panel installers to get more accurate cost estimates and further refine their decision-making process.

4) Cost-benefit analysis of installing a heat pump:

User: What is the monetary advantage of installing a heat pump?

EC: To determine the monetary advantage of installing a heat pump, please provide the following information:

- Climate: What is the local climate in your area? The efficiency and potential savings of a heat pump depend on the temperature range it will operate in.
- Current heating and cooling system: What type of heating and cooling system do you currently have in your home? This will help determine the potential energy and cost savings from switching to a heat pump.
- Energy prices: What are the current prices for electricity and any other energy sources (e.g., natural gas, oil, propane) used in your home? This will be used to compare the operating costs of different systems.
- Heat pump efficiency: What is the efficiency rating (e.g., Seasonal Energy Efficiency Ratio or SEER, and Heating Seasonal Performance Factor or HSPF) of the heat pump you are considering? Higher efficiency ratings indicate a more energy-efficient system.
- Installation and maintenance costs: What are the estimated costs of purchasing, installing, and maintaining the heat pump? These costs will factor into the overall monetary advantage of the system.

Once you provide this information, we can help you estimate the potential monetary advantages of installing a heat pump.

User: (note that the user directly answers the questions)

- Mild
- Central air conditioning
- 0.75 per kWh
- I don't know the rating
- 200 per year

Background processing:

$$\text{minimize } C \in \mathbb{R} \quad C \quad (4a)$$

$$\text{subject to } C = a - b, \quad (4b)$$

$$b \geq 0, \quad (4c)$$

$$a \geq 0, \quad (4d)$$

$$b = c \times 0.75, \quad (4e)$$

$$a = d \times 0.75 + 200, \quad (4f)$$

where C , a , b , c , d represent total cost, heat pump cost, central AC cost, central AC energy consumption and heat pump energy consumption.

```
# Code (only given with important lines)
central_AC_cost = cp.Variable()
heat_pump_cost = cp.Variable()
constraints = [
    central_AC_cost >= 0,
    heat_pump_cost >= 0,
    central_AC_cost == central_AC_energy_consumption
    * 0.75,
    heat_pump_cost == heat_pump_energy_consumption *
    0.75 + 200,
]
problem = cp.Problem(cp.Minimize(heat_pump_cost -
    central_AC_cost), constraints)
problem.solve()
# Output:
Annual savings with heat pump: $550.00
#end
```

EC: With the calculated annual savings of \$550.00, you can now use this information to make an informed decision about whether or not to

install a heat pump in your home. Here's how you can use this information:

- Evaluate the investment payback period: Compare the upfront cost of purchasing and installing the heat pump to the annual savings. By dividing the total upfront cost by the annual savings, you can determine the number of years it will take for the heat pump to pay for itself through energy savings. For example, if the heat pump costs \$5,000 to purchase and install, the payback period would be about 9.1 years ($\$5,000 / \550).
- Compare with other options: Use the annual savings to compare the heat pump with alternative heating and cooling systems. Research the energy efficiency, operating costs, and maintenance costs of other systems (such as furnaces, boilers, or air source heat pumps) to determine if the heat pump is the most cost-effective and environmentally friendly option for your needs.
- Assess your long-term plans: Consider how long you plan to stay in your current home. If the payback period is longer than the time you plan to live in the house, it may not make financial sense to install the heat pump. However, if you expect to live in the house for a long time, the heat pump could provide significant long-term savings.
- Evaluate environmental impact: By calculating the difference in energy consumption between the central air conditioning system and the heat pump, you can estimate the environmental impact of each system. Installing a heat pump can help reduce your home's carbon footprint, contributing to a more sustainable future.
- Apply for incentives or rebates: Some governments or utilities offer financial incentives for installing energy-efficient heating and cooling systems like heat pumps. Research local and national rebate programs to see if you qualify for any financial assistance, which could further reduce the upfront cost and shorten the payback period.
- Consult with a professional: Before making a decision, consult with a heating, ventilation, and air conditioning (HVAC) professional who can assess your home's specific needs, recommend the best heat pump model, and provide accurate cost and savings estimates.

By analyzing this information and considering these factors, you can make an educated decision about whether installing a heat pump is the right choice for your home.

5) *Battery sizing optimization for long-term electricity bill reduction:* Figuring out the optimal sizing of a home battery system to minimize long-term electricity costs calls for a balance between the initial capital outlay and anticipated savings, taking into account variables like demand and the price of electricity [31]. Here, our objective is not to arrive at a professionally precise computation but to provide insight to users about the diverse factors playing a role in such decisions.

User: I live in a sunny location with some sunlight throughout the day. However, I cannot sell excess energy to the grid. I want to buy a battery to sustain my electricity usage in the evenings which is 30 kWh every day. Price of the batteries increases with the battery sizes' square. I want to learn what would be the best battery size to reduce my total cost including the battery cost and buying electricity from the grid.

EC:

- Battery Cost: Could you provide a function or a method to calculate the cost of the battery based on its capacity? As you mentioned that the price of batteries increases with the battery size's square, we'd need the specifics to come up with a calculation.
- Efficiency of the battery system: This is important because not every kilowatt-hour (kWh) of electricity that you put into the battery will be available for use – some will be lost in the charging and discharging process.
- Amount of Solar Generation: How much solar energy (in kWh) are you able to generate per day? This will be critical to understand how much excess energy you can store in your battery system during the day.
- Electricity Rate: How much are you currently paying per kWh for electricity from the grid? This will help in comparing the cost of grid electricity with the cost of storing and using solar energy.
- Battery Lifespan: How many years does the battery last? This will be important for calculating the overall cost over the battery's lifespan, allowing us to work out a cost per kWh over the battery's lifetime.

User: (note that the user directly answers the questions)

- Battery cost is the price of the battery times the square of battery capacity, where the price of the battery is 10 dollars per kilowatt hour.
- 0.95
- It produces 10kwh each day
- \$0.25
- 2 years

Background processing:


$$\text{minimize}_{B \in \mathbb{R}} \quad P_{\text{bat}} B^2 + (P_e \times 365 \times Y)G \quad (5a)$$

$$\text{subject to} \quad G \geq D - S - B * \eta_{\text{bat}} \quad (5b)$$

$$G \geq 0, \quad (5c)$$

where P_{bat} , B , P_e , Y , G , D , S , η_{bat} are unit cost of the battery, battery size, price of electricity from the grid, number of years to consider, electricity usage from grid each day, demand in the evenings, solar energy production in the morning and the efficiency of the battery, respectively.

```
# Code (only given with important lines)
battery_capacity = cp.Variable(1)
battery_cost = price_battery * battery_capacity**2
grid_electricity_cost = electricity_rate * cp.maximum(
    (30 - battery_efficiency * battery_capacity -
     solar_generation), 0)
total_cost = battery_cost + 365 * battery_lifespan *
    grid_electricity_cost
problem = cp.Problem(cp.Minimize(total_cost))
problem.solve()
# Output:
Optimal Battery Size: 9.12505953102464 kWh
```

 **EC:** The optimal battery size of approximately 9.12 kWh means that this is the most cost-effective battery size for your specific situation, based on the parameters you provided.

Here's how you can use this information:

- Purchasing a battery: When shopping for a home battery storage system, look for one that offers a usable capacity of around 9.12 kWh. If the exact size isn't available, you should consider the ones closest to this capacity.
- Understanding your energy use: This calculation shows that a 9.12 kWh battery should store enough solar energy during the day to cover your typical electricity usage in the evenings.
- Maximizing your savings: By using a 9.12 kWh battery, you should be able to minimize your reliance on the grid, thereby reducing your electricity bills over the two-year lifespan of the battery.
- Evaluating your options: While this analysis provides a theoretical optimum, other practical considerations may apply when selecting a battery system. You may want to consider other factors such as warranty period, brand reputation, physical size, maintenance requirements, and potential for future expansion of your solar system.

Firstly, EC accurately identifies the objective, encompassing the initial purchase and installation costs and the yearly electricity expenses over the chosen period. It also successfully encodes the appropriate energy balance. Notably, in this instance, all the details offered by the users are wholly utilized in the formulation. One potential limitation is that the optimal battery size is lower than the total solar production, which could lead to under-utilization of some solar energy. However, this would require a more granular, hourly formulation to address, which might not always have the necessary data available. The explanations that EC provides are thorough and balanced, and should aid users in their decision-making.