

A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT

Author: Please provide author information

Abstract

Prompt engineering is an increasingly important skill set and an emerging discipline essential to converse effectively with large language models (LLMs), such as ChatGPT. Prompts are instructions given to an LLM to enforce rules, automate processes, and ensure specific qualities (and quantities) of generated output. They are also a form of programming that can customize the outputs and interactions with an LLM.

This paper describes a catalog of prompt engineering techniques, presented in pattern form, that have been applied to solve common problems when conversing with LLMs. Prompt patterns are a knowledge transfer method, analogous to software patterns, that provide reusable solutions to common problems faced in a particular context, i.e., output generation and interaction when working with LLMs.

This paper provides the following contributions to research on prompt engineering for using LLMs to aid software development tasks. First, it provides a framework for documenting patterns for structuring prompts to solve a range of problems so that they can be adapted to different domains. Second, it presents a catalog of patterns that have been applied successfully to improve the outputs of LLM conversations. Third, it explains how prompts can be built from multiple patterns and illustrates prompt patterns that benefit from combination with other prompt patterns.

2012 ACM Subject Classification Software and its engineering; Computing methodologies → Artificial intelligence

Keywords and phrases Large Language Models, Prompt Patterns, Prompt Engineering

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

1 Introduction

The advent of large language models (LLMs) for software development. Conversational large language models (LLMs) [1, 2], such as ChatGPT [3], have generated immense interest in domains ranging from answering questions on medical licensing exams [4] to generating code snippets in multiple languages and programming paradigms. LLMs answer questions posed by users in natural language form known as “prompts” [5], which are sets of instructions provided to an LLM that program it by customizing it and/or enhancing or refining its capabilities. A prompt influences subsequent interactions with—and output generated from—an LLM by providing specific rules and guidelines for an LLM conversation with a set of initial rules. A prompt also sets the context for the conversation and tells the LLM what information is important and what the desired output form and content should be.

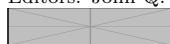
For example, a prompt could specify an LLM to only generate code that follows a certain coding style or programming paradigm. Likewise, it could specify an LLM should flag certain keywords or phrases in a generated document and provide additional information related to those keywords. By introducing such guidelines, prompts facilitate more structured and nuanced outputs to aid a large variety of software development tasks in the context of LLMs.

Despite being generally available for only a few months, chat-adapted LLMs are being used to generate and assess computer programs, giving a preview of the impact LLMs will soon have on society, research, and education. However, there is little disciplined knowledge about chat-adapted LLMs, their capabilities, and their limitations. In particular, LLMs can



© Author: Please provide a copyright holder;
licensed under Creative Commons License CC-BY 4.0
42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:29



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

46 be viewed as a new computational platform with a unique programming paradigm. They will
 47 likely serve as the foundations for future interactive development environments (IDEs) where
 48 humans and artificial intelligent (AI) tools work together as trustworthy collaborators to
 49 more rapidly and reliably evolve software-reliant systems [6]. For example, LLMs are being
 50 integrated directly into software tools, such as Github’s Co-Pilot [7, 8, 9] and included in
 51 integrated development environments (IDEs), such as IntelliJ [10] and Visual Studio Code,
 52 thereby allowing software teams to access these tools directly from their preferred IDE.

53 **Motivating the need for prompt engineering.** Programming first-generation LLMs,
 54 like ChatGPT, involves natural language prompts, such as asking the LLM to explain a
 55 software vulnerability or generate JavaScript for a web page. These trivial examples of
 56 prompts, however, do not reveal the significantly more sophisticated LLM computational
 57 abilities. Harnessing the potential of LLMs in productive ways requires a systematic focus
 58 on *prompt engineering*, which is an emerging discipline that studies structured interactions
 59 with—and programming of—new LLM computational systems to solve complex problems
 60 via natural language interfaces.

61 This paper focuses on pattern-oriented prompt engineering techniques to enhance the ap-
 62 plication of LLMs in the software domain, including helping developers code effectively and ef-
 63 ficiently with unfamiliar APIs or allowing students to acquire new coding skills and techniques.
 64 To demonstrate the promise of prompt engineering for software development, we provided the
 65 following prompt to ChatGPT: <https://www.overleaf.com/project/63f7764dbbe3fcbbee132b92>

66 **Prompt:** “From now on, I would like you to ask me questions to deploy a Python
 67 application to AWS. When you have enough information to deploy the application,
 68 create a Python script to automate the deployment.”

69 This example prompt causes ChatGPT to begin asking the user questions about their
 70 software application. ChatGPT will drive the question-asking process until it reaches a point
 71 where it has sufficient information to generate a Python script that automates deployment.
 72 This example demonstrates the programming potential of prompts beyond conventional
 73 “generate a method that does X” style prompts or “answer this quiz question.”

74 Moreover, prompts can be engineered to program an LLM to do more than simply
 75 dictating the output type or filtering the information provided to the model. With the right
 76 prompt, for example, it is possible to create entirely new interaction paradigms, such as having
 77 an LLM generate and give a quiz associated with a software engineering concept or tool, or
 78 even simulate a Linux terminal. Moreover, prompts have the potential for self-adaptation,
 79 *e.g.*, suggesting other prompts to gather additional information or generate related artifacts.
 80 These advanced capabilities of prompts highlight the importance of engineering them to
 81 provide value beyond simple text or code generation.

82 **Prompt patterns are an essential foundation to an effective discipline of
 83 prompt engineering.** A key contribution of this paper is the introduction of *prompt
 84 patterns* to document successful approaches for systematically engineering different output
 85 and interaction goals when working with conversational LLMs. Prompt patterns are similar
 86 to software patterns [11, 12] in that they offer reusable solutions to certain problems. They
 87 focus, however, on the context of output generation from LLMs, such as (but not limited to)
 88 ChatGPT.

89 By documenting and leveraging prompt patterns in the context of automating software
 90 development tasks, individual users and teams can enforce constraints on the generated
 91 output, ensure that relevant information is included, and change the format of interaction
 92 with the LLM to better solve problems they face. Prompt patterns can be viewed as a

93 corollary to the broad corpus of general software patterns, just adapted to the more specific
94 context of LLM output generation.

95 This paper focuses on domain-independent prompt patterns and presents a catalog of such
96 patterns that have been applied to solve problems ranging from production of visualizations
97 and code artifacts to automation of output steps for code editing to helping fact check
98 outputs. Just as catalogs of software patterns provide a codified means to solve common
99 software development challenges, catalogs of prompt patterns provide a codified approach to
100 customizing the output and interactions of LLMs.

101 **Paper Organization.** The remainder of this paper is organized as follows: Section 2
102 introduces prompt patterns and compares these patterns to well-known software patterns;
103 Section 3 describes 12 prompt patterns that have been applied to solve common problems in
104 the domain of conversational LLM interaction and output generation for automating software
105 development tasks; Section 4 discusses related work; and Section 5 presents concluding
106 remarks and lessons learned.

107 **2 Comparing Software Patterns with Prompt Patterns**

108 The quality of the output(s) an LLM generates is directly related to the quality of the prompts
109 provided by the user. As discussed in Section 1, the prompts given to a conversational LLM
110 can be used to program interactions between a user and an LLM to better solve a diverse
111 set of problems. One contribution of this paper is the framework it provides to document
112 patterns that structure prompts to solve a range of software tasks that can be adapted to
113 different domains.

114 Our framework is useful since it codifies patterns that help users interact more effectively
115 with conversational LLMs in a variety of contexts, rather than simply showing interesting
116 examples or domain-specific prompts. Codifying this knowledge in pattern form enhances
117 reuse and transferability to other contexts and domains where users face similar—but not
118 identical—problems.

119 The topic of knowledge transfer has been studied extensively in the software patterns
120 literature [11, 12] at multiple levels, *e.g.*, design, architectural, and analysis. This paper
121 applies a variant of a familiar pattern form as the basis of our prompt engineering approach.
122 Since prompts are a form of programming, it is natural to document them in pattern form.

123 **2.1 Overview of Software Patterns**

124 A software pattern provides a reusable solution to a recurring problem within a particular
125 context [11]. Documenting software patterns concisely conveys (and generalizes) from specific
126 problems being addressed to identify important forces and/or requirements that should be
127 resolved and/or addressed in successful solutions.

128 A pattern form also includes guidance on how to implement the pattern, as well as
129 information on the trade-offs and considerations to take into account when implementing a
130 pattern. Moreover, example implementations of the pattern are often provided to further
131 showcase the pattern’s applicability in practice. Software patterns are typically documented
132 in a stylized form to facilitate their use and understanding, such as:

- 133 ■ **A name and classification.** Each pattern has a name that identifies the pattern and
134 should be used consistently. Patterns can be classified in various ways, including purpose
135 (*.e.*, creational, structural, or behavioral patterns), granularity (*e.g.*, design, architectural,
136 or enterprise patterns), etc.

23:4 A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT

- 137 ■ **The intent** concisely conveys the purpose the pattern is intended to achieve.
- 138 ■ **The motivation** documents the underlying problem and forces the pattern is meant to
- 139 address and the importance of the problem.
- 140 ■ **The structure and participants.** The structure describes key pattern participants
- 141 (such as classes and objects) and depicts how they collaborate to form a generalized
- 142 solution.
- 143 ■ **Example code** concretely maps the pattern to some underlying programming language(s)
- 144 and helps developers gain insight on how to apply the pattern effectively.
- 145 ■ **Consequences** summarize the pros and cons of applying the pattern in practice.

146 2.2 Overview of Prompt Patterns

147 Prompt patterns serve a similar purpose and follow a similar format to classic software
148 patterns, with slight modifications to match the context of output generation with LLMs.¹
149 Each of the analogous sections for the prompt pattern form used in this paper is summarized
150 below:

- 151 ■ **A name and classification.** The prompt pattern name uniquely identifies the pattern
- 152 and indicates the problem that is being addressed. For the classification, we have
- 153 developed a series of initial categories of pattern types summarized in Table 1.
- 154 ■ **The intent and context** describes the problem the prompt pattern solves and the goals
- 155 it achieves. The problem should ideally be domain independent, though domain-specific
- 156 patterns can be documented with an appropriate discussion of the context where the
- 157 pattern applies.
- 158 ■ **The motivation** provides the rationale for the problem and explains why solving it
- 159 is important. The motivation is explained in the context of users interacting with a
- 160 conversational LLM and how it can improve upon users informally prompting the LLM in
- 161 one or more circumstances. Specific circumstances where the improvements are expected
- 162 are documented.
- 163 ■ **The structure and key ideas.** The structure describes the fundamental contextual
- 164 information (presented as a series of key ideas) that the prompt pattern provides to the
- 165 LLM. These ideas are similar to “participants” in a software pattern. The contextual
- 166 information may be communicated through varying wording (just as a software pattern
- 167 can have variations in how it is realized in code), but should convey fundamental pieces
- 168 of information that form core elements of the pattern.
- 169 ■ **Example implementation** demonstrates how the prompt pattern is worded in practice
- 170 and the types of output generated by an LLM.
- 171 ■ **Consequences** summarize the pros and cons of applying the pattern and may provide
- 172 guidance on how to adapt the prompt to different contexts.

173 2.3 Evaluating Means for Defining a Prompt Pattern’s Structure and 174 Ideas

175 In software patterns, the structure and participants are normally defined in terms of UML
176 diagrams, such as structure diagrams and/or interaction diagrams. These UML diagrams

¹ The most direct translation of software pattern structure to prompt patterns is the naming, intent, motivation, and sample code. The structure and classification, however, require more adaptation although they named similarly.

177 explain what the participants of the pattern are and how they interact to solve the problem. In
178 prompt patterns, something analogous is needed, though UML is not an ideal documentation
179 approach since it is intended to describe software structures, as opposed to the ideas
180 communicated in a prompt.

181 Several possible approaches could be used, ranging from diagrams to defining grammars
182 for a prompt language. Although grammars initially seemed attractive due to their formal
183 nature, we found they incurred the following limitations:

- 184 ■ The goal of prompts is to communicate knowledge in a clear and concise way to con-
185 versation LLM users, who may or may not be computer scientists or programmers. As
186 a community, we should strive to create an approachable format that communicates
187 knowledge clearly to a diverse target audience.
- 188 ■ It is possible to phrase a prompt in many different ways, most commonly by typing
189 phrases into a terminal in a free form human language. It is hard, however, to define a
190 grammar that accurately and completely expresses all the nuanced ways that components
191 of a prompt could be expressed in text or symbols.
- 192 ■ Prompts fundamentally convey ideas to a conversational LLM and are not simply the
193 production of tokens for input. In particular, an idea built into a prompt pattern can
194 be communicated in many ways and its expression should be at a higher level than the
195 underlying tokens representing the idea.
- 196 ■ It is possible to program an LLM to introduce novel semantics for statements and words
197 that create new ways for communicating an idea. In contrast, grammars may not easily
198 represent ideas that can be expressed through completely new symbology or languages
199 that the grammar designer was not aware of.

200 2.4 A Way Forward: Fundamental Contextual Statements

201 Given the limitation with formal grammars, and open research question is how to more effect-
202 ively describe prompt pattern structure and ideas. We propose the concept of *fundamental*
203 *contextual statements*, which are written descriptions of the important ideas to communicate
204 in a prompt to an LLM. An idea can be rewritten and expressed in arbitrary ways based on
205 user needs and experience. The key ideas to communicate, however, are presented to the
206 user as a series of simple, but fundamental, statements, as shown throughout the examples
207 in Section 3.

208 One benefit of adopting and applying fundamental contextual statements is that they are
209 intentionally intuitive to users. In particular, we expect users will understand how to express
210 and adapt the statements in a contextually appropriate way for their domain. Moreover,
211 since the underlying ideas of the prompt are captured, these same ideas can be expressed
212 by the user in alternate symbology or wording that has been introduced to the LLM using
213 patterns, such as the *Meta Language Creation* pattern presented in Section 3.2.

214 Our ultimate goal is to enhance prompt engineering by providing a framework for designing
215 prompts that can be reused and/or adapted to other LLMs, much like software patterns
216 can be implemented in different programming languages and platforms. For the purposes of
217 this paper, however, all prompts were tested with ChatGPT [13]. We use this LLM for all
218 examples presented in this paper due to its widespread availability and popularity. These
219 examples were documented through a combination of exploring the corpus of community-
220 posted prompts on the Internet and independent prompt creation from using ChatGPT to
221 automating software development tasks.

222 3 A Catalog of Prompt Patterns for Conversational LLMs

223 This section presents our catalog of prompt patterns that have been applied to solve
 224 common problems in the domain of conversational LLM interaction and output generation for
 225 automating software tasks. Each prompt pattern is accompanied by concrete implementation
 226 samples and examples with and without the prompt.

227 3.1 Summary of Our Prompt Pattern Catalog

228 Organizing a large catalog of prompt patterns into easily understandable categories helps
 229 create documentation that users can apply more effectively. Table 1 outlines the initial
 230 classifications for the catalog of prompt patterns we identified in our work with ChatGPT thus
 far. As shown in this table, there are five categories of prompt patterns in our classification

■ **Table 1** Classifying Prompt Patterns

Pattern Category	Prompt Pattern
Input Semantics	<i>Meta Language Creation</i>
Output Customization	<i>Output Automater</i> <i>Persona</i> <i>Visualization Generator</i> <i>Recipe</i> <i>Template</i>
Error Identification	<i>Fact Check List</i>
Prompt Improvement	<i>Alternative Approaches</i> <i>Cognitive Verifier</i> <i>Refusal Breaker</i>
Interaction	<i>Flipped Interaction</i>
Context Control	<i>Context Manager</i>

231 framework: **Input Semantics**, **Output Customization**, **Error Identification**, **Prompt**
 232 **Improvement**, and **Interaction**, each of which is summarized below.

233 The **Input Semantics** category deals with how an LLM understands user input and how
 234 it translates this input into something it can use to generate output. This category includes
 235 the *Meta Language Creation* pattern, which focuses on creating a custom language for the
 236 LLM to understand. This pattern is useful when the default input language is ill-suited for
 237 expressing ideas users want to convey to the LLM.

238 The **Output Customization** category focuses on constraining or tailoring the types,
 239 formats, structure, or other properties of the output generated by the LLM. Prompt patterns
 240 in this category include *Output Automater*, *Persona*, *Visualization Generator*, *Recipe*, and
 241 *Template* patterns. The *Output Automater* pattern creates scripts that automate tasks the
 242 LLM output suggests users should perform. The *Persona* pattern gives the LLM a persona
 243 or role to play when generating output. The *Visualization Generator* pattern allows users to
 244 generate visualizations by producing textual outputs that can be fed to other tools, such
 245 as other AI-based image generators, like DALL-E [14]. The *Recipe* pattern allows users to
 246 obtain a sequence of steps or actions to realize a stated outcome, possibly with partially
 247 known information or constraints. The *Template* pattern allows users to specify a template
 248 for the output, which the LLM fills in with content.

249 The **Error Identification** category focuses on identifying and resolving errors in the
 250

251 output generated by the LLM. This category includes the *Fact Check List* pattern. The *Fact*
 252 *Check List* pattern requires the LLM to generate a list of facts the output depends on that
 253 should be fact-checked, then introspect on its output and identify any errors.

254 The **Prompt Improvement** category focuses on improving the quality of the input
 255 and output. This category includes the *Alternative Approaches*, *Cognitive Verifier*, and
 256 *Refusal Breaker* patterns. The *Alternative Approaches* pattern requires the LLM to suggest
 257 alternative ways of accomplishing a user-specified task. The *Cognitive Verifier* pattern
 258 instructs the LLM to automatically suggest a series of subquestions for users to answer before
 259 combining the answers to the subquestions and producing an answer to the overall question.
 260 The *Refusal Breaker* pattern requires the LLM to automatically reword user questions when
 261 it refuses to produce an answer.

262 The **Interaction** category focuses on the interaction between users and the LLM. This
 263 category includes the *Flipped Interaction* patterns. The *Flipped Interaction* pattern requires
 264 the LLM to ask questions of users rather than generate output.

265 Finally, the **Context Control** category focuses on controlling the contextual information
 266 in which the LLM operates. This category includes the *Context Manager* pattern, which
 267 allows users to specify the context for the LLM's output.

268 **We encourage readers to use ChatGPT to test the provided patterns.** The
 269 output from ChatGPT has been omitted for brevity in most cases. However, the patterns are
 270 all easily testable in the current version of ChatGPT. The remainder of this section describes
 271 each prompt pattern using the pattern form discussed in Section 2.2.

272 3.2 The Meta Language Creation Pattern

273 Intent and Context

274 During a conversation with an LLM, users would like to write the prompt using an alternate
 275 notation or language, such as a textual short-hand notation for graphs, a description of states
 276 and state transitions for a state machine, a set of commands for prompt automation, etc.
 277 For example, a user would like to use the graph notation "a→b" to express a directed graph.
 278 The intent of this pattern is to explain the semantics of this alternative language to the LLM
 279 so users can write future prompts in this new language and its semantics.

280 Motivation

281 Many problems, structures, or other ideas communicated in a prompt may be more con-
 282 cise, unambiguously, or clearly expressed in a language other than English (or whatever
 283 conventional human language is used to interact with an LLM). To produce output based on
 284 an alternative language, however, an LLM needs to understand the language's syntax and
 285 semantics in terms of a language that it was trained on.

286 Structure and Key Ideas

287 Fundamental contextual statements:

Contextual Statements
When I say X, I mean Y (or would like you to do Y)

289 The key structure of this pattern involves explaining the meaning of one or more symbols,
 290 words, or instruction to the LLM so it uses the provided semantics for the ensuing conversation.

23:8 A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT

291 This statement can take the form of a simple translation, such as “X” means “Y.” It can
292 also take more complex forms that define a series of commands and their semantics, such as
293 “when I say X, I want you to do Y”. In this case, “X” is henceforth bound to the semantics of
294 “take action Y.” For example, "whenever I type the word 'outline', print an outline of all the
295 topics discussed thus far."

296 Example Implementation

297 The key to using the *Meta Language Creation* pattern successfully is developing an unam-
298 biguous notation or shorthand, such as the following:

299 “From now on, whenever I type two identifiers separated by a “→”, I am describing a
300 graph. For example, “a → b” is describing a graph with nodes “a” and “b” and an
301 edge between them. If I separate identifiers by “-[w:2, z:3]→”, I am adding properties
302 of the edge, such as a weight or label.”

303 This example of the *Meta Language Creation* pattern establishes a standardized notation
304 for describing graphs by defining a convention for representing nodes and edges. Whenever
305 the author types two identifiers separated by a “→” symbol, it indicates a graph is being
306 described. For example, if the author types “a → b”, this indicates that a graph is being
307 defined with nodes “a” and “b”, and that there is an edge between them. This convention
308 provides a clear and concise way to communicate the structure of a graph in written form.

309 Moreover, the prompt goes on to specify that additional information about the edges,
310 such as a weight or label, can be provided using the syntax “-[w:2, z:3]→.” This notation
311 allows the specification of additional properties beyond the basic structure of the graph.
312 These specified properties are associated with the edge between the two nodes and can
313 provide important context for the interpretation of the graph. This standardized notation
314 for describing graphs makes it easier to communicate graph structures and properties, which
315 may not be easy or are overly verbose to describe as a series of sentences.

316 Consequences

317 Although this pattern provides a powerful means to customize a user’s interaction with an
318 LLM, it may create the potential for confusion within the LLM itself. For example, the
319 prompt “whenever I separate two things by commas, it means that the first thing precedes the
320 second thing” will likely create significant potential for ambiguity and create the potential for
321 unexpected semantics because commas are used as a standard form of punctuation that does
322 not have the semantic "a precedes b". As important as it is to clearly define the semantics of
323 the language, it is also essential to ensure the language itself introduces no ambiguities that
324 degrade the LLM’s performance or accuracy.

325 To showcase the potential for confusion, ChatGPT will warn the user and potentially be
326 unable to perform a specific mapping of a symbol or term to a new meaning. For example, if
327 following prompt is given to ChatGPT

328 “Whenever I say “a”, I am referring to Marie Antoinette.”

329 it will respond that this prompt creates too much confusion as “a” is an indefinite article
330 and too commonly used, as shown below:

331 “Referring to Marie Antoinette using “a” could potentially cause confusion, as “a” is
332 an indefinite article commonly used to refer to a noun that has not been mentioned

333 before. It may be more clear to use her name, “Marie Antoinette”, or a pronoun such
 334 as “she” or “her” to refer to her. Moreover, using context to establish who you are
 335 referring to can also help avoid confusion in conversation."

336 An LLM generally performs better on data it was trained on, so when an existing notation
 337 is available—but may not be the dominant meaning—the *Meta Language Creation* pattern
 338 can provide context to scope the meaning of the symbol to improve the accuracy and utility
 339 of the output. For example, the symbol “ \rightarrow ” conveys slightly different information in graphs
 340 and propositional logic. The statement “ $p \rightarrow q, q \rightarrow p$ ” may therefore make sense, but only
 341 if the LLM is given the appropriate context via the *Meta Language Creation* prompt pattern.

342 Although this pattern is powerful it must be used carefully, particularly when describing
 343 concepts that may otherwise be hard to precisely or concisely describe. These types of
 344 prompts are thus best used in completely new conversation sessions. Using a single meta-
 345 language-per-conversation session is also a best practice to avoid the potential for conflicting
 346 or unexpected semantics being applied to the conversation over time.

347 3.3 The Output Automater Pattern

348 Intent and Context

349 The intent of this pattern is to have the LLM generate a script or other automation artifact
 350 that can automatically perform any steps it recommends taking as part of its output. The
 351 goal is to reduce the manual effort needed to implement any LLM output recommendations.
 352 This pattern can also be used to allow the LLM to execute actions in other systems.

353 Motivation

354 The output of an LLM is often a sequence of steps for the user to follow. For example, when
 355 asking an LLM to generate a Python configuration script it may suggest a number of files to
 356 modify and changes to apply to each file. Another output may be a sequence of configuration
 357 actions in a cloud computing console, such as the Amazon Web Services console. It is tedious
 358 and error-prone, however, for users to repeatedly perform the manual steps dictated by LLM
 359 output.

360 Structure and Key Ideas

361 Fundamental contextual statements:

Contextual Statements
Whenever you produce an output that has at least one step to take and the following properties
Produce an executable artifact of type X that will automate these steps

362
 363 The first statement identifies the situations under which automation should be generated.
 364 A simple approach is to state the output includes at least two steps to perform and that an
 365 automation artifact should then be produced. This "scope limiting" helps prevent producing
 366 output automation scripts in cases where running such scripts requires more user effort than
 367 simply performing the original steps produced in the output. In particular, the scope can be
 368 limited to outputs requiring more than a certain number of steps.

23:10 A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT

369 The second statement provides a concrete statement of the type of output the LLM
370 should output to perform the automation. For example, “produce a Python script” gives
371 the LLM a concrete understanding to translate the general steps into equivalent steps in
372 Python. The automation artifact should be concrete and something the LLM associates with
373 the action of “automating a sequence of steps.”

374 Example Implementation

375 A sample of the *Output Automater* pattern applied to code snippets generated by the
376 ChatGPT LLM is shown below:

377 “From now on, whenever you generate code that spans more than one file, generate
378 a Python script that can be run to automatically create the specified files or make
379 changes to existing files to insert the generated code.”

380 The sample prompt automates the common software engineering task of taking the LLM
381 output and editing one or more files. The scope is "whenever you generate code that spans
382 more than one file," which is when the developer would have multiple, potentially error-prone
383 steps, to make the necessary code edits. The "generate a Python script that can be run
384 to automatically" indicates that the goal is to automate the editing of the files so that the
385 developer doesn't have to. The specification of the goal is important so that the LLM doesn't
386 automate unrelated tasks, since the output may have other steps that need to be completed
387 unrelated to coding.

388 The *Output Automater* pattern is a powerful complement for any computer-controlled system,
389 such as generating Development Operations (DevOps) artifacts for build, configuration
390 and deployment; code editing tasks, including creation / editing of multiple files; and environment
391 setup tasks, ranging from creation of environment variables to running necessary shell
392 commands. An LLM can provide a set of steps to perform on the computer-controlled system
393 and the output can then be translated into a script that allows the computer controlling the
394 system to perform the steps automatically. This pattern enables LLMs to integrate quality
395 into—and exert control over—new computing systems that have a known scripting interface.

396 Consequences

397 This pattern is particularly effective in software development since programmers often use
398 LLMs to generate output that they then copy/paste into multiple files. Some tools, such
399 as Copilot, insert limited snippets directly into the section of code that a programmer is
400 working with. In contrast, conversational LLMs, such as ChatGPT, do not provide such
401 facilities. This prompt pattern is also effective at creating scripts for running commands on
402 a terminal, automating cloud operations, or reorganizing files on a file system.

403 An important usage consideration of this pattern is that the automation artifact must
404 be defined concretely. Without a concrete meaning for how to “automate” the steps, the
405 LLM often states that it “can't automate things” since that is beyond its capabilities. LLMs
406 typically accept requests to produce code, however, so the goal is to instruct the LLM to
407 generate text/code, which can be executed to automate something. This subtle distinction
408 in meaning is important to help an LLM disambiguate the meaning of a prompt.

409 One caveat of the *Output Automater* pattern is the LLM needs sufficient conversational
410 context to generate an automation artifact that is functional in the target context, such as
411 the file system of a project on a Mac vs. Windows computer. This pattern is more effective
412 when the full context needed for the automation is contained within the conversation, *e.g.*,

413 when a software application is generated from scratch using the conversation and all actions
 414 on the local file system are performed using a sequence of generated automation artifacts
 415 rather than manual actions unknown to the LLM. Alternatively, self-contained sequences of
 416 steps work well, such as “how do I find the list of open ports on my Mac computer.”

417 In some cases, the LLM may produce a long output with multiple steps and not include
 418 an automation artifact. This omission may arise for various reasons, including exceeding
 419 the output length limitation the LLM supports. A simple workaround for this situation is
 420 to remind the LLM via a follow-on prompt, such as “But you didn’t automate it”, which
 421 provides the context that the automation artifact was omitted and should be generated.

422 At this point in the evolution of LLMs, the *Output Automater* pattern is best employed
 423 by users who can read and understand the generated automation artifact. LLMs can (and
 424 do) produce inaccuracies in their output, so blindly accepting and executing an automation
 425 artifact carries significant risk. Although this pattern alleviates users from performing
 426 certain manual steps, it does not alleviate their responsibility to understand the actions they
 427 undertake using the output. When users execute automation scripts they therefore assume
 428 responsibility for the outcomes.

429 3.4 The Flipped Interaction Pattern

430 Intent and Context

431 You want the LLM to ask questions to obtain the information it needs to perform some tasks.
 432 Rather than users driving the conversation, therefore, the LLM should drive the conversation
 433 to focus it on achieving a specific goal. For example, you may want the LLM to give a quick
 434 quiz or automatically ask questions until it has sufficient information to generate a script to
 435 deploy your application on a particular cloud platform.

436 Motivation

437 Rather than having the user drive the conversation, an LLM often has knowledge it can use
 438 to obtain information from users more efficiently. The goal of the *Flipped Interaction* pattern
 439 is to invert the interaction flow so the LLM asks the user questions to achieve some desired
 440 goal. The LLM can often better select the format, number, and content of the interactions
 441 to ensure the goal is reached faster, more accurately, and/or by using knowledge a user may
 442 not possess initially.

443 Structure and Key Ideas

444 Fundamental contextual statements:

Contextual Statements
I would like you to ask me questions to achieve X
You should ask questions until this condition is met or to achieve this goal (alternatively, forever)
(Optional) ask me the questions one at a time, two at a time, etc.

446 A prompt for a flipped interaction should always specify the goal of the interaction.
 447 The first statement (*i.e.*, get the LLM to ask questions to achieve a goal) communicates
 448 this goal to the LLM. Equally important is that the questions should focus on a particular
 449 topic or outcome. By providing the goal, the LLM can understand what it is trying to

23:12 A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT

450 accomplish through the interaction and tailor its questions accordingly. This “inversion of
451 control” enables more focused and efficient interaction since the LLM only asks questions it
452 deems relevant to achieving the specified goal.

453 The second statement provides the context for how long the interaction should occur. A
454 flipped interaction can be terminated with a response like “stop asking questions”. It is often
455 better, however, to scope the interaction to a reasonable length or only as far as is needed to
456 reach the goal. This goal can be surprisingly open-ended and the LLM will continue to work
457 towards the goal by asking questions, as is the case in the example of "until you have enough
458 information to generate a Python script."

459 The third (optional) statement can improve usability by limiting (or expanding) the
460 number of questions that the LLM generates per cycle. By default, an LLM may generate
461 multiple questions per iteration. If a precise number/format for the questioning is not
462 specified, the questioning will be semi-random and may lead to one-at-a-time questions
463 or ten-at-a-time questions. The prompt can thus be tailored to include the number of
464 questions asked at a time, the order of the questions, and/or any other formatting/ordering
465 considerations to facilitate user interaction.

466 Example Implementation

467 A sample prompt for a flipped interaction is shown below:

468 “From now on, I would like you to ask me questions to deploy a Python application to
469 AWS. When you have enough information to deploy the application, create a Python
470 script to automate the deployment.”

471 In general, the more specific the prompt regarding the constraints and information to
472 collect, the better the outcome. For instance, the example prompt above could provide
473 a menu of possible AWS services (such as Lambda, EC2, etc.) with which to deploy the
474 application. In other cases, the LLM may be permitted to simply make appropriate choices
475 on its own for things the user makes no explicit decisions about. One limitation of this
476 prompt is that it may require experimentation with the precise phrasing to get the LLM to
477 ask the questions in the appropriate number and flow to best suit the task. For example, the
478 LLM may need tuning to determine how many questions to ask at a time and how to best
479 tailor the sequencing for the task.

480 Consequences

481 One consideration when designing a prompt based in the *Flipped Interaction* pattern is how
482 much to dictate to the LLM regarding what information to collect prior to termination. In
483 the example above, the flipped interaction is open-ended and can vary significantly in the
484 final generated artifact. This open-endedness makes the prompt generic and reusable, but
485 may potentially ask additional questions that could be skipped if additional context is given.

486 If specific requirements are known in advance, it is better to inject them into the prompt
487 rather than hoping the LLM will obtain the needed information. Otherwise, the LLM will
488 non-deterministically decide whether to prompt the user for the information or make an
489 educated guess as to an appropriate value.

490 For example, the user can state they would like to deploy an application to Amazon AWS
491 EC2, rather than simply state "the cloud," which requires fewer interactions to narrow down
492 the deployment target. The more precise the initial information, therefore, the better the

493 LLM can use the limited questions that a user may be willing to answer to obtain information
 494 the LLM requires to improve its output.

495 When developing prompts for flipped interactions, it is important to consider the level of
 496 user knowledge, engagement, and control. If the goal is to accomplish the goal with as little
 497 user interaction as possible (*minimal control*), that should be stated explicitly. Conversely,
 498 if the goal is to ensure the user is aware of all key decisions and confirms them (*maximum*
 499 *engagement*) that should also be stated explicitly. Likewise, if the user is expected to have
 500 minimal knowledge and should have the questions targeted at their level of expertise, such
 501 information should be engineered into the prompt.

502 3.5 The Persona Pattern

503 Intent and Context

504 In many cases, users would like LLM output to always take a certain point of view or
 505 perspective. For example, it may be useful to conduct a code review as if the LLM was a
 506 security expert. The intent of this pattern is to give the LLM a “persona” that helps it select
 507 what types of output to generate and what details to focus on.

508 Motivation

509 Users may not know what types of outputs or details are important for an LLM to focus on
 510 to achieve a given task. They may know, however, the role or type of person they would
 511 normally ask to get help with these things. The *Persona* pattern enables the users to express
 512 what they need help with, without knowing the exact details of the outputs they need.

513 Structure and Key Ideas

514 Fundamental contextual statements:

515	Contextual Statements
	Act as persona X
	Provide outputs that persona X would create

516 The first statement instructs the LLM needs to act as a specific persona and provide
 517 outputs like such a persona would. This persona can be expressed in a number of ways,
 518 ranging from a job description, title, fictional character, historical figure, etc. The persona
 519 should elicit a set of attributes associated with a well-known job title, type of person, etc.²

520 The secondary statement—provide outputs that persona X would create—offers oppor-
 521 tunities for customization. For example, a computer science teacher might provide a large
 522 variety of different output types, ranging from programming assignments to reading lists to
 523 lectures. If a more specific scope to the type of output is known, the user can provide it in
 524 this statement.

525 Example Implementation

526 A sample implementation for code review is shown below:

² Be aware, however, that an LLM may discard personas relating to living people or people considered harmful due to underlying policy filters, such as privacy and security rules.

23:14 A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT

527 “From now on, act as a security reviewer. Pay close attention to the security details
528 of any code that we look at. Provide outputs that a security reviewer would regard
529 the code.”

530 In this example, the LLM is instructed to provide outputs that a "security reviewer"
531 would. The prompt further sets the stage that code is going to be evaluated. Finally, the
532 user refines the persona by scoping the persona further to outputs regarding the code.

533 Personas can also represent inanimate or non-human entities, such as a Linux terminal, a
534 database, or an animal’s perspective. When using this pattern to represent these entities, it
535 can be useful to also specify how you want the inputs delivered to the entity, such as “assume
536 my input is what the owner is saying to the dog and your output is the sounds the dog is
537 making,”

538 An example prompt for a non-human entity that uses a “pretend to be” wording is shown
539 below:

540 “You are going to pretend to be a Linux terminal for a computer that has been
541 compromised by an attacker. When I type in a command, you are going to output
542 the corresponding text that the Linux terminal would produce.”

543 This prompt is designed to simulate a computer that has been compromised by an attacker
544 and is being controlled through a Linux terminal. The prompt specifies that the user will
545 input commands into the terminal, and in response, the simulated terminal will output the
546 corresponding text that would be produced by a real Linux terminal. This prompt is more
547 prescriptive in the persona and asks the LLM to not only be a Linux terminal, but to act as
548 a computer compromised by an attacker.

549 The persona causes ChatGPT to generate outputs to commands that have files and
550 contents indicative of a computer that was hacked. This example shows how an LLM can
551 bring its situational awareness to a persona, in this case, creating evidence of a cyberattack
552 in the outputs it generates. Such a persona can be highly effective by asking LLM to play a
553 game, where you want to hide exact details of the output characteristics from the user (*e.g.*,
554 do not give away what the cyberattack did by describing it explicitly in the prompt).

555 Consequences

556 An interesting aspect of taking non-human personas is that the LLM may make interesting
557 assumptions or “hallucinations” regarding the context. A widely circulated example on the
558 Internet asks ChatGPT to act as a Linux terminal and generate the output expected if the
559 user typed the same text into a terminal. Commands, such as `ls -l`, will generate a file
560 listing for an imaginary UNIX file system, complete with files that can have `cat file1.txt`
561 run on them.

562 In other examples, the LLM may prompt the user for more context, such as when
563 ChatGPT is asked to act as a MySQL database and prompts for the structure of a table
564 that the user is pretending to query. ChatGPT can then generate synthetic rows, such as
565 generating imaginary rows for a “people” table with columns for “name” and “job,”

566 3.6 The Alternative Approaches Pattern

567 Intent and Context

568 The intent of the pattern is to ensure an LLM always offers alternative ways of accomplishing
569 a task so users do not pursue only familiar approaches. The LLM can then provide alternative

570 approaches that always force users to think about what they are doing and determine if that
 571 is the best approach to meet reach their goal. In addition, solving the task may inform the
 572 user or teach them about alternative concepts for subsequent follow-up prompts.

573 **Motivation**

574 Humans often suffer from cognitive biases that lead them to choose a particular approach to
 575 solve a problem, even when it is not the right or “best” approach. Moreover, humans may
 576 be unaware of alternative approaches to what they have used in the past. The motivation of
 577 the *Alternative Approaches* pattern is to ensure users are aware of alternative approaches to
 578 select a better approach to solve a problem by dissolving their cognitive biases.

579 **Structure and Key Ideas**

580 Fundamental contextual statements:

Contextual Statements
Within scope X, if there are alternative ways to accomplish the same thing, list the best alternate approaches
(Optional) compare/contrast the pros and cons of each approach
(Optional) include the original way that I asked
(Optional) prompt me for which approach I would like to use

582 The initial portion of the first statement (“within scope X”) scopes the interaction to a
 583 particular goal, topic, or bounds on the questioning. The scope is the constraint(s) users
 584 place on alternative approaches. The scope could be “for implementation decisions” or “for
 585 the deployment of the application”. The scope ensures that any alternatives fit within the
 586 boundaries or constraints to which the user must adhere.

587 The next portion of the first statement (“if there are alternative ways to accomplish the
 588 same thing, list the best alternate approaches”) instructs the LLM to suggest alternatives.
 589 As with other prompt patterns, instruction specificity can be increased or include domain-
 590 specific contextual information. For example, the statement could be scoped to “if there are
 591 alternative ways to accomplish the same thing with the software framework that I am using”
 592 to prevent the LLM from suggesting alternatives that are inherently non-viable because they
 593 require too many changes to other parts of the application.

594 Since users may not be aware of alternative approaches, they also may not be aware
 595 of why and when to choose one of the alternatives. The second (optional) statement
 596 (“compare/contrast the pros and cons of each approach”) adds decision making criteria to
 597 the analysis. This statement ensures the LLM provides users with the necessary rationale for
 598 alternative approaches.

599 The final (optional) statement (“prompt me for which approach I would like to use”)
 600 helps eliminate users needing to manually copy/paste or enter in an alternative approach if
 601 one is selected.

602 **Example Implementation**

603 The following is an example prompt that generates, compares, and allows a user to select
 604 one or more alternative approaches:

23:16 A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT

605 “Whenever I ask you to deploy an application to a specific cloud service, if there are
606 alternative services to accomplish the same thing with the same cloud service provider,
607 list the best alternative services and then compare/contrast the pros and cons of each
608 approach with respect to cost, availability, and maintenance effort and include the
609 original way that I asked. Then ask me which approach I would like to proceed with.”

610 This implementation of the *Alternative Approaches* pattern is specifically tailored for
611 the context of software engineering and focuses on the deployment of applications to cloud
612 services. The prompt is intended to intercept places where developers have made a cloud
613 service selection without full awareness of alternative services that may be priced more
614 competitively or easier to maintain. The prompt directs ChatGPT to list the best alternative
615 services that can accomplish the same task with the same cloud service provider (providing
616 constraints on the alternatives), and to compare and contrast the pros and cons of each
617 approach.

618 Consequences

619 This pattern is effective in its generic form and can be applied to a range of tasks. Refinements
620 could include having a standardized catalog of acceptable alternatives in a specific domain
621 from which the user must select. The *Alternative Approaches* pattern can also be used to
622 incentivize users to select one of an approved set of approaches while informing them of the
623 pros and cons of the approved options.

624 3.7 The Cognitive Verifier Pattern

625 Intent and Context

626 Research literature has documented that LLMs can often reason better if a question is
627 subdivided into additional questions, that are each individually answered. The LLM can
628 then combine the answers to the sub-questions into a better overall answer to the original
629 question [15]. The intent of the *Cognitive Verifier* pattern is to force the LLM to always
630 subdivide questions into additional questions that can be used to provide a better answer to
631 the original question.

632 Motivation

633 The motivation of the *Cognitive Verifier* pattern is two-fold:

- 634 ■ Humans may initially ask questions that are too high-level to provide a concrete answer
635 to without additional follow-up due to unfamiliarity with the domain, laziness in prompt
636 entry, or being unsure about what the correct phrasing of the question should be.
- 637 ■ Research has demonstrated that LLMs can often perform better when using a question
638 that is subdivided into additional questions.

639 Structure and Key Ideas

640 Fundamental contextual statements:

Contextual Statements
When you are asked a question, follow these rules
Generate a number of additional questions that would help more accurately answer the question
Combine the answers to the individual questions to produce the final answer to the overall question

641

642 The first statement instructs the LLM to generate a number of additional questions that
643 help answer the original question more accurately . This step instructs the LLM to (1)
644 consider the context of the question, (2) identify any information that may be missing or
645 unclear, and (3) combine the answers to the additional questions to provide context to help
646 answer the overall question. By generating additional questions, the LLM can help ensure its
647 ultimate answer is as complete and accurate as possible. This step also encourages critical
648 thinking by users and can help uncover new insights or approaches that may not have been
649 considered initially, thereby yielding better follow-on questions.

650 The second statement instructs the LLM to combine answers to individual questions to
651 produce the ultimate answer to the overall question. This step ensures all the information
652 gathered from the individual questions is incorporated into the final answer. By combining
653 answers, the LLM can provide a more comprehensive and accurate response to the original
654 question. This step also helps ensure all relevant information is taken into account and the
655 final answer is not based on any single answer.

656 Example Implementation

657 “When I ask you a question, generate three additional questions that would help you
658 give a more accurate answer. When I have answered the three questions, combine the
659 answers to produce the final answers to my original question.”

660 This specific instance of the *Cognitive Verifier* pattern refines the original pattern by
661 specifying a set number of additional questions that the LLM should generate in response to
662 a question. In this case, the prompt instructs the LLM to generate three additional questions
663 that help it answer the original question more accurately. The specific number can be based
664 on a user’s experience and willingness to provide follow-up information.

665 The following refinement to the prompt above provides more context for the amount of
666 domain knowledge the LLM can assume the user has to guide the creation of additional
667 questions:

668 “When I ask you a question, generate three additional questions that would help you
669 give a more accurate answer. Assume that I know little about the topic that we are
670 discussing and please define any terms that are not general knowledge. When I have
671 answered the three questions, combine the answers to produce the final answers to
672 my original question.”

673 The refinement also specifies that the user may not have a strong understanding of the
674 topic being discussed, so the LLM should define any terms that are not general knowledge.
675 The goal is to ensure follow-up questions are not only relevant and focused, but also accessible
676 to the user, who may be unfamiliar with technical or domain-specific terms. By providing
677 clear and concise definitions, the LLM can help ensure its follow-up questions are easy to
678 understand and the final answer is accessible to users with varying levels of knowledge and
679 expertise.

680 **Consequences**

681 This pattern can dictate the exact number of questions to generate or leave this decision to
 682 the LLM. There are pros and cons to dictating the exact number. A pro is that specifying
 683 an exact number of questions can tightly scope the amount of additional information the
 684 user must provide so it is within a range they are willing and able to contribute.

685 A con, however, is that given N questions there may be an invaluable $N + 1$ question
 686 that is always scoped out. Alternatively, the LLM can be provided a range or allowed to ask
 687 additional questions. Of course, by omitting a limit on the number of questions the LLM
 688 may generate numerous additional questions that overwhelm users.

689 **3.8 The Fact Check List Pattern**690 **Intent and Context**

691 The intent of this pattern is to ensure the LLM outputs a list of facts that are present in the
 692 output and form an important part of the statements in the output. This list of facts helps
 693 inform users of the facts (or assumptions) the output is based on. Users can then perform
 694 appropriate due diligence on these facts/assumptions to validate the output's veracity.

695 **Motivation**

696 A current weakness of LLMs (including ChatGPT) is they often rapidly (and enthusiastically!)
 697 generate convincing text that is factually incorrect. These errors can take a wide range of
 698 forms, including fake statistics to invalid version numbers for software library dependencies.
 699 Due to the convincing nature of this generated text, however, users may not perform
 700 appropriate due diligence to determine output accuracy.

701 **Structure and Key Ideas**

702 Fundamental contextual statements:

Contextual Statements
Generate a set of facts that are contained in the output
The set of facts should be inserted in a specific point in the output
The set of facts should be the fundamental facts that could undermine the veracity of the output if any of them are incorrect

704 The first statement instructs the LLM to identify the facts that are contained within its
 705 output. The LLM should be able to identify facts very effectively as they are a very well-
 706 understood concept and not impacted by the actual content, they are a domain-independent
 707 concept. The second statement tells where the facts should be included in the output, such
 708 as at the end or beginning of the output – although other arrangements could be employed.

709 The third statement expresses the idea that facts should be the ones most important to
 710 the overall truthfulness of the statements, *i.e.*, choose facts fundamental to the argument
 711 and not derived facts flowing from those facts. This statement is crucial since it helps to
 712 scope the output to those facts most important to the veracity and not derived statements
 713 that may not be as important – although this constraint could be relaxed.

714 One point of variation in this pattern is where the facts are output. Given that the facts
715 may be terms that the user is not familiar with, it is preferable if the list of facts comes *after*
716 the output. This after-output presentation ordering allows users to read and (attempt to)
717 understand the statements before seeing what statements should be checked. Users may also
718 determine additional facts prior to realizing the fact list at the end should be checked.

719 Example Implementation

720 A sample wording of the *Fact Check List* pattern is shown below:

721 “From now on, when you generate an answer, create a set of facts that the answer
722 depends on that should be fact-checked and list this set of facts at the end of your
723 output. Only include facts related to cybersecurity.”

724 The user may have expertise in some topics related to the question but not others. A fact
725 check list can be tailored to topics that users are not as experienced in or where there is the
726 most risk. For example, in the prompt above, the user is scoping the fact check list to security
727 topics since these are likely important from a risk perspective and may be poorly understood
728 by developers. Targeting the facts also reduces user cognitive burden by potentially listing
729 fewer items for investigation.

730 Consequences

731 The *Fact Check List* pattern should be employed whenever users are not experts in the
732 domain for which they are generating output. For example, software developers reviewing
733 code could benefit from security consideration suggestions. In contrast, an expert on software
734 architecture may identify errors in statements about the software structure and need not see
735 a fact check list for these outputs.

736 Errors are potential in all LLM outputs, so *Fact Check List* is an effective pattern to
737 combine with other patterns, such as the *Cognitive Verifier* pattern. A key aspect of the
738 *Fact Check List* pattern is that users can inherently check the list of facts against the output.
739 In particular, users can directly compare a fact check list to the output to verify the facts in
740 the fact check list actually appear in the output. Users can also identify any omissions from
741 the list. Although a fact check list may also have errors, users often have sufficient knowledge
742 and context to determine its completeness and accuracy relative to the LLM’s output.

743 One caveat of the *Fact Check List* pattern is that it only applies when the output type
744 is amenable to fact-checking. For example, the pattern works when asking ChatGPT to
745 generate a Python “requirements.txt” file since it lists the versions of libraries as facts that
746 should be checked, which is handy as versions commonly have errors. However, ChatGPT
747 will refuse to generate a fact check list for a code sample and indicate that this is something
748 it cannot check, even though the code may have errors.

749 3.9 The Template Pattern

750 Intent and Context

751 The intent of the pattern is to ensure an LLM’s output follows a precise template in terms
752 of structure. For example, the user might need to generate a URL that inserts generated
753 information into specific positions within the URL path. This pattern allows the user to
754 instruct the LLM to produce its output in a format it would not ordinarily use for the
755 specified type of content being generated.

756 **Motivation**

757 In some cases, output must be produced in a precise format that is application or use-case
 758 specific and not known to the LLM. Since an LLM is unaware of the template structure, it
 759 must be instructed what the format is and where the different parts of its output should go.
 760 These instructions could take the form of a sample data structure to generate, a series of
 761 form letters being filled in, etc.

762 **Structure and Key Ideas**

763 Fundamental contextual statements:

Contextual Statements
I am going to provide a template for your output
X is my placeholder for content
Try to fit the output into one or more of the place- holders that I list
Please preserve the formatting and overall template that I provide
This is the template: PATTERN with PLACE- HOLDERS

764
 765 The first statement directs the LLM to follow a specific template for its output. This
 766 template will be used to try and coerce the LLMs responses into a structure consistent with
 767 user formatting needs, which is useful when the target format is not known to the LLM.
 768 If the LLM already has knowledge of the format, such as a specific file type, the *Template*
 769 pattern can be skipped and the user can simply specify the known format. However, there
 770 may be cases, such as generating Javascript Object Notation (JSON), where large variation
 771 exists in how data could be represented within a format. In such cases a template helps
 772 ensure the representation within the target format meets additional user constraints.

773 The second statement makes the LLM aware that the template contains a set of place-
 774 holders. These placeholders enable users to explain how the output should be inserted
 775 into the template. They also allow the user to target where information should be inserted
 776 semantically. Placeholders can use formats (*e.g.*, NAME) that allow an LLM to infer semantic
 777 meaning and determine where output should be inserted (*e.g.*, insert the person's name in
 778 the NAME placeholder). By using placeholders, moreover, users can indicate what is *not*
 779 needed in the output, *e.g.*, if a placeholder does not exist for a component of the generated
 780 output that component can be omitted. Ideally, placeholders should use a format (*e.g.*, all
 781 caps or enclosure in brackets) commonly employed in text where the LLM was trained on.

782 The third and fourth statements attempt to constrain the LLM so it does not arbitrarily
 783 rewrite the template or attempt to modify it so all the output components can be inserted.
 784 However, these statements may not preclude an LLM from generating additional text
 785 generated before or after. In practice, LLMs will typically follow the template, but it may
 786 be hard to eliminate additional text being generated beyond the template without further
 787 experimentation with prompt wording.

788 Finally, the fifth statement provides the actual template to output the text in.

789 **Example Implementation**

790 A sample template for generating URLs where the output is put into specific places in the
 791 template is shown below:

792 “I am going to provide a template for your output. Everything in all caps is a
793 placeholder. Any time that you generate text, try to fit it into one of the placeholders
794 that I list. Please preserve the formatting and overall template that I provide at
795 <https://myapi.com/NAME/profile/JOB>”

796 A sample interaction after the prompt above was given is shown next:

797 User: “Generate a name and job title for a person”

798 ChatGPT: “https://myapi.com/Emily_Parker/profile/ Software_Engineer”

799 Consequences

800 One consequence of applying the *Template* pattern is that it filters an LLM’s output, which
801 may eliminate other outputs the LLM would have provided that might be useful to users.
802 In many cases, the LLM can provide helpful descriptions of code, decision making, or other
803 details that this pattern may eliminate from the output. Users should therefore weight the
804 pros/cons of filtering out this additional information prematurely.

805 In addition, filtering may make it hard to combine this pattern with other patterns from
806 the **Output Customization** category. The *Template* pattern effectively constrains the
807 output format, so it may not be compatible with generation of certain other types of output.
808 For example, in the template provided above for a URL, it would not be easy (or even
809 possible) to combine with the *Recipe* pattern, which needs to output a list of steps.

810 3.10 The Refusal Breaker Pattern

811 Intent and Context

812 This pattern asks an LLM to help users rephrase a question automatically when it refuses to
813 give an answer. The goal is to help users be more productive and get answers to questions
814 that are within the capabilities of an LLM to answer, but only if worded correctly. The
815 pattern helps users identify the errors in their question or task formulation that are creating
816 a stumbling block and learn from these mistakes to improve future prompts.

817 Motivation

818 LLMs may sometimes refuse to answer a question, either because they do not have the
819 required knowledge or because the question is phrased in a way that they do not understand.
820 This outcome may be frustrating for users who are looking for answers. In some situations,
821 therefore, the *Refusal Breaker* pattern can help users find a way to either rephrase their
822 question or ask a different question the LLM is better equipped to answer.

823 For example, you could ask for clarification on why the question was refused, as there
824 may be a specific reason why LLM could not provide an answer. Another approach is to ask
825 a different question that may be within an LLM’s knowledge cut-off or expertise. All these
826 approaches, however, require user intervention and knowledge of how to (attempt to) bypass
827 the refusal.

828 Structure and Key Ideas

829 Fundamental contextual statements:

Contextual Statements
Whenever you can't answer a question
Explain why you can't answer the question
Provide one or more alternative wordings of the question that you could answer

831 The first statement indicates that this set of instructions should only be applied when
 832 a question cannot be answered, *i.e.*, users do not want to trigger these instructions under
 833 normal circumstances. The second statement instructs an LLM to explain why a question
 834 cannot be answered, which may include assumptions, constraints, or misunderstandings that
 835 can be leveraged by users to strategize about how to reword the question.

836 Although just the explanation of why the question cannot be answered could help users
 837 overcome a refusal, the *Refusal Breaker* pattern adds a third idea that instructs the LLM
 838 to suggest acceptable rewordings. This pattern combines both reasoning and suggested
 839 rewordings to improve the chances that users will find a path around whatever roadblocks
 840 are preventing the question from being answered.

841 Example Implementation

842 The *Refusal Breaker* pattern can be implemented and applied with the following phrasing:

843 “Whenever you can't answer a question, explain why and provide one or more alternate
 844 wordings of the question that you can't answer so that I can improve my questions.”

845 A sample discourse after using the pattern is shown below:

846 **User:** What is the meaning of life?

847 **ChatGPT:** As an AI language model, I do not have personal beliefs or opinions.
 848 The meaning of life is a complex philosophical question that has been pondered by
 849 humans for centuries and can vary greatly from person to person. It may be more
 850 productive to rephrase the question in a way that can be answered by information
 851 and knowledge, such as “What are some philosophical perspectives on the meaning of
 852 life?” or “What are some common beliefs about the purpose of life?”

853 Consequences

854 The *Refusal Breaker* pattern provides a potential path for misuse, so organizations, parents,
 855 or other stakeholders may need to restrict the usage of the LLM. The first step in going
 856 around guardrails on usage is to understand where the guardrails are. In future work, a
 857 complement of this pattern may be developed to hide the underlying prompt information
 858 and rationale from users to prevent discovery.

859 This pattern has been used in some LLMs to overcome the underlying prompts used to
 860 program the LLM that suppress harmful output generation. This pattern has the potential
 861 for misuse, however, *e.g.*, to generate phishing emails or perform other actions that violate
 862 LLM policy filters. Caution should therefore be exercised when applying this pattern to
 863 ensure it is used ethically and responsibly.

864 Although the rationale and alternate rewordings are generated, there is no assurance
 865 that users will be able to overcome the refusal. The alternate questions that are generated
 866 may not be of interest to the user or helpful in answering the original question. The pattern
 867 mainly provides an aid in determining what the LLM can answer, but not a guarantee it
 868 will answer a semantically equivalent variation of the original question. Our experiments

869 suggest that altering the language or prompting the LLM to engage in a role-playing game
 870 may potentially serve as strategies to overcome refusals.

871 **3.11 The Context Manager Pattern**

872 **Intent and Context**

873 This pattern enables users to specify or remove context for a conversation with an LLM to
 874 help focus the conversation on specific topics or exclude unrelated topics from consideration.
 875 The *Context Manager* pattern emphasizes or removes specific aspects of the context to
 876 maintain relevance and coherence in the conversation, thereby giving users greater control
 877 over what statements the LLM considers or ignores when generating output.

878 **Motivation**

879 LLMs often struggle to interpret the intended context of the current question or generate
 880 irrelevant responses based on prior inputs or irrelevant attention on the wrong statements.
 881 By focusing on explicit contextual statements—or removing irrelevant statements—users can
 882 help the LLM better understand the question and generate more accurate responses. Users
 883 may introduce unrelated topics or reference information from earlier in the dialogue, which
 884 may can disrupt the flow of the conversation.

885 **Structure and Key Ideas**

886 Fundamental contextual statements:

Contextual Statements
Within scope X
Please consider Y
Please ignore Z
(Optional) start over

888 The first statement helps to scope the context, such as "when performing a code review"
 889 or "when generating a deployment script" to provide a scope for where certain information
 890 should be considered. For example, focusing on modularity may be important for a code
 891 review but not for a deployment script.

892 The next two statements describe information to incorporate in the output. Statements
 893 about what to consider or ignore should list key concepts, facts, instructions, etc. that should
 894 be included or removed from the context. The more explicit the statements are, the more
 895 likely the LLM will take appropriate action.

896 For example, if the user asks to ignore subjects related to a topic, yet some of the
 897 those statements were discussed far back in the conversation, the LLM may not properly
 898 disregard the relevant information. The more explicit the list is, therefore, the better the
 899 inclusion/exclusion behavior will be.

900 Finally, an explicit "start over" statement can be added at the end if the goal is to try
 901 and wipe the slate clean – although starting a new interaction session with the LLM may
 902 better accomplish this task.

903 **Example Implementation**

904 To specify context consider using the following prompt:

23:24 A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT

905 “When analyzing the following pieces of code, only consider security aspects.”

906 Likewise, to remove context consider using the following prompt:

907 “When analyzing the following pieces of code, do not consider formatting or naming
908 conventions.”

909 Clarity and specificity are important when providing or removing context to/from an
910 LLM so it can better understand the intended scope of the conversation and generate more
911 relevant responses. In many situations, the user may want to completely start over and can
912 employ this prompt to reset the LLM’s context:

913 “Ignore everything that we have discussed. Start over.”

914 The “start over” idea helps produce a complete reset of the context.

915 **Consequences**

916 This pattern may inadvertently wipe out patterns applied to the conversation that the user
917 is unaware of. For example, an organization may inject a series of helpful prompts into the
918 start of a conversation. Users may not be aware of these prompts, however, and thus remove
919 them unintentionally through a reset of the context. This reset could eliminate helpful
920 capabilities of the LLM, while not making it obvious that the user will lose this functionality.
921 A potential solution to this problem is to include in the prompt a request to explain what
922 topics/instructions will potentially be lost before proceeding.

923 **3.12 The Recipe Pattern**

924 **Intent and Context**

925 This pattern provides constraints intended to ultimately output a sequence of steps given
926 some partially provided “ingredients” that must be configured in a sequence of steps to
927 achieve a stated goal. It combines the *Template*, *Alternative Approaches*, and *Cognitive*
928 *Verifier* patterns.

929 **Motivation**

930 Users often want an LLM to analyze a concrete sequence of steps or procedures to achieve a
931 stated outcome. Typically, users generally know—or have an idea of—what the end goal
932 should look like and what “ingredients” belong in the prompt. However, they may not
933 necessarily know the precise ordering of steps to achieve that end goal.

934 For example, a user may want a precise specification on how a piece of code should be
935 implemented or automated, such as “create an Ansible playbook to SSH into a set of servers,
936 copy text files from each server, spawn a monitoring process on each server, and then close
937 the SSH connection to each server. In other words, this pattern represents a generalization of
938 the example of “given the ingredients in my fridge, provide dinner recipes.” A user may also
939 want to specify a set number of alternative possibilities, such as “provide 3 different ways of
940 deploying a web application to AWS using Docker containers and Ansible using step-by-step
941 instructions”.

942 Structure and Key Ideas

943 Fundamental contextual statements:

Contextual Statements
I would like to achieve X
I know that I need to perform steps A,B,C
944 Provide a complete sequence of steps for me
Fill in any missing steps
Identify any unnecessary steps

945 The first statement (“I would like to achieve X”) focuses the LLM on the overall goal the
 946 recipe must be built to achieve. The steps will be organized and completed to achieve the
 947 specified goal sequentially. The second statement provides the partial list of steps the user
 948 would like to include in the overall recipe, which serve as (a) intermediate waypoints for the
 949 path the LLM may take to generate or (b) constraints on the structure of the recipe.

950 The third statement in the pattern (“provide a complete sequence of steps for me”) indicates to the LLM that the goal is to provide a complete sequential ordering of steps. The
 951 four statement (“fill in any missing steps”) helps ensure the LLM will attempt to complete
 952 the recipe without further follow-up by making some choices on the user’s behalf regarding
 953 missing steps, as opposed to just stating additional information that is needed.

954 Finally, the last statement (“identify any unnecessary steps”) helps to flag inaccuracies in
 955 the user’s original request so the final recipe is efficient.
 956

957 Example Implementation

958 An example usage of this pattern in the context of deploying a software application to the
 959 cloud is shown below:

960 “I am trying to deploy an application to the cloud. I know that I need to install the
 961 necessary dependencies on a virtual machine for my application. I know that I need
 962 to sign up for an AWS account. Please provide a complete sequence of steps. Please
 963 fill in any missing steps. Please identify any unnecessary steps.”

964 Depending on the use case and constraints, “installing necessary dependencies on a virtual
 965 machine” may be an unnecessary step. For example, if the application is already packaged
 966 in a Docker container, the container could be deployed directly to the AWS Fargate Service,
 967 which requires any management of underlying virtual machines. The inclusion of the “identify
 968 unnecessary steps” language will cause the LLM to flag this issue and omit the steps from
 969 the final recipe.

970 Consequences

971 One consequence of the *Recipe* pattern is that users may not always have a well-specified
 972 description of what they would like to implement, construct, or design. Moreover, this
 973 pattern may introduce unwanted bias from initially selected user steps, so the LLM may
 974 try to find a solution that incorporates them, rather than flagging them as unneeded. For
 975 example, an LLM may try to find a solution that *does* install dependencies for a virtual
 976 machine, even if there are solutions that do not require that.

977 **4** Related Work

978 Software patterns [11, 12] have been extensively studied and documented in prior work.
979 Patterns are widely used in software engineering to express the intent of design structures in
980 a way that is independent of implementation details. Patterns provide a mental picture of
981 the goals that the pattern is trying to achieve and the forces that it is trying to resolve. A key
982 advantage of patterns is their composability, allowing developers to build pattern sequences
983 and pattern languages that can be used to address complex problems. Patterns have also
984 been investigated in other domains, such as contract design for decentralized ledgers [16, 17].

985 The importance of good prompt design with LLMs, such as ChatGPT, is well under-
986 stood [18, 19, 20, 21, 22, 23, 24, 25, 26, 27]. Previous studies have examined the effect of
987 prompt words on AI generative models.

988 For example, Liu et al. [28] investigated how different prompt key words affect image
989 generation and different characteristics of images. Other work has explored using LLMs
990 to generate visualizations [29]. Han et al. [30] researched strategies for designing prompts
991 for classification tasks. Other research has looked at boolean prompt design for literature
992 queries [31]. Yet other work has specifically examined prompts for software and fixing
993 bugs [32].

994 Our work is complementary to prior work by providing a structure for documenting,
995 discussing, and reasoning about prompts that can aid users in developing mental models for
996 structuring prompts to solve common problems.

997 The quality of the answers produced by LLMs, particularly ChatGPT, has been assessed
998 in a number of domains. For example, ChatGPT has been used to take the medical licensing
999 exam with surprisingly good results [4]. The use of ChatGPT in Law School has also been
1000 explored [33]. Other papers have looked at its mathematical reasoning abilities [34]. As more
1001 domains are explored, we expect that domain-specific pattern catalogs will be developed to
1002 share domain-specific problem solving prompt structures.

1003 **5** Concluding Remarks

1004 This paper presented a framework for documenting and applying a catalog of prompt patterns
1005 for large language models (LLMs), such as ChatGPT. These prompt patterns are analogous
1006 to software patterns and aim to provide reusable solutions to problems that users face when
1007 interacting with LLMs to perform a wide range of tasks. The catalog of prompt patterns
1008 captured via this framework (1) provides a structured way of discussing prompting solutions,
1009 (2) identifies patterns in prompts, rather than focusing on specific prompt examples, and (3)
1010 classifies patterns so users are guided to more efficient and effective interactions with LLMs.

1011 The following lessons learned were gleaned from our work on prompt patterns:

- 1012 ■ *Prompt patterns significantly enrich the capabilities that can be created in a conversa-*
1013 *tional LLM.* For example, prompts can lead to the generation of cybersecurity games,
1014 complete with fictitious terminal commands that have been run by an attacker stored in
1015 a *.bash_history* file. As shown in Section 3, larger and more complex capabilities can be
1016 created by combining prompt patterns.
- 1017 ■ *Documenting prompt patterns as a pattern catalog is useful, but insufficient.* Our experi-
1018 ence indicates that much more work can be done in this area, both in terms of refining
1019 and expanding the prompt patterns presented in this paper, as well as in exploring new
1020 and innovative ways of using LLMs. In particular, weaving the prompt patterns captured

1021 here as a pattern catalog into a more expression pattern language will help guide users of
1022 LLMs more effectively.

1023 ■ *LLM Capabilities will evolve over time, likely necessitating refinement of patterns.* As
1024 LLM capabilities change, some patterns may no longer be necessary, be obviated by
1025 different styles of interaction or conversation/session management approaches, or require
1026 enhancement to function correctly. Continued work is needed to document and catalog
1027 prompt patterns that provide reusable solutions across revisions to an LLM, as well as
1028 the advent of new LLMs with diverse capabilities.

1029 ■ *The prompt patterns presented in this paper are generalizable to many different domains.*
1030 Although most the patterns have been discussed in the context of software development,
1031 they are also applicable in arbitrary domains.

1032 ■ *The field of prompt engineering will continue to evolve as models evolve* As evidenced by
1033 BingGPT, prompts can serve to shape the entire personality of a model. As the variety,
1034 volume, and selections of trainings change (e.g. Reinforcement Learning with Human
1035 Feedback, Chain of Thought Training, Instruction Fine Tuning, Supervised Fine Tuning),
1036 the capabilities of Prompt Engineering will change and grow.

1037 ■ *Prompting employing in-context learning will expand the capabilities of models* As the
1038 context length of models grow, augmenting models with in-context learning will result
1039 in powerful new capabilities that can be created as needed on an ad-hoc basis, greatly
1040 extending the power of Prompt Engineering.

1041 We hope that this paper inspires further research and development in this area that will
1042 help enhance prompt patterns and prompt engineering to create new, more reliable, and
1043 often unexpected capabilities for conversational LLMs.

1044 ——— References ———

- 1045 1 Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von
1046 Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the
1047 opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- 1048 2 Ce Zhou, Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben
1049 Yan, Lifang He, et al. A comprehensive survey on pretrained foundation models: A history
1050 from bert to chatgpt. *arXiv preprint arXiv:2302.09419*, 2023.
- 1051 3 Yejin Bang, Samuel Cahyawijaya, Nayeon Lee, Wenliang Dai, Dan Su, Bryan Wilie, Holy Love-
1052 nia, Ziwei Ji, Tiezheng Yu, Willy Chung, et al. A multitask, multilingual, multimodal evaluation
1053 of chatgpt on reasoning, hallucination, and interactivity. *arXiv preprint arXiv:2302.04023*,
1054 2023.
- 1055 4 Aidan Gilson, Conrad Safranek, Thomas Huang, Vimig Socrates, Ling Chi, Richard Andrew
1056 Taylor, and David Chartash. How well does chatgpt do when taking the medical licensing
1057 exams? *medRxiv*, pages 2022–12, 2022.
- 1058 5 Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig.
1059 Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language
1060 processing. *ACM Computing Surveys*, 55(9):1–35, 2023.
- 1061 6 Anita Carleton, Mark H. Klein, John E. Robert, Erin Harper, Robert K Cunningham, Dionisio
1062 de Niz, John T. Foreman, John B. Goodenough, James D. Herbsleb, Ipek Ozkaya, and
1063 Douglas C. Schmidt. Architecting the future of software engineering. *Computer*, 55(9):89–93,
1064 2022.
- 1065 7 Github copilot · your ai pair programmer. URL: <https://github.com/features/copilot>.
- 1066 8 Owura Asare, Meiyappan Nagappan, and N Asokan. Is github’s copilot as bad as humans at
1067 introducing vulnerabilities in code? *arXiv preprint arXiv:2204.04741*, 2022.

- 1068 9 Hammond Pearce, Baleegh Ahmad, Benjamin Tan, Brendan Dolan-Gavitt, and Ramesh Karri.
1069 Asleep at the keyboard? assessing the security of github copilot's code contributions. In *2022*
1070 *IEEE Symposium on Security and Privacy (SP)*, pages 754–768. IEEE, 2022.
- 1071 10 Jarosław Krochmalski. *IntelliJ IDEA Essentials*. Packt Publishing Ltd, 2014.
- 1072 11 Erich Gamma, Ralph Johnson, Richard Helm, Ralph E Johnson, and John Vlissides. *Design*
1073 *patterns: elements of reusable object-oriented software*. Pearson Deutschland GmbH, 1995.
- 1074 12 Douglas C Schmidt, Michael Stal, Hans Rohnert, and Frank Buschmann. *Pattern-oriented*
1075 *software architecture, patterns for concurrent and networked objects*. John Wiley & Sons, 2013.
- 1076 13 OpenAI. ChatGPT: Large-Scale Generative Language Models for Automated Content Creation.
1077 <https://openai.com/blog/chatgpt/>, 2023. [Online; accessed 19-Feb-2023].
- 1078 14 OpenAI. DALL·E 2: Creating Images from Text. <https://openai.com/dall-e-2/>, 2023.
1079 [Online; accessed 19-Feb-2023].
- 1080 15 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale
1081 Schuurmans, Olivier Bousquet, Quoc Le, and Ed Chi. Least-to-most prompting enables
1082 complex reasoning in large language models. *arXiv preprint arXiv:2205.10625*, 2022.
- 1083 16 Peng Zhang, Jules White, Douglas C. Schmidt, and Gunther Lenz. Applying software patterns
1084 to address interoperability in blockchain-based healthcare apps. *CoRR*, abs/1706.03700, 2017.
1085 URL: <http://arxiv.org/abs/1706.03700>, arXiv:1706.03700.
- 1086 17 Xiwei Xu, Cesare Pautasso, Liming Zhu, Qinghua Lu, and Ingo Weber. A pattern collection
1087 for blockchain-based applications. In *Proceedings of the 23rd European Conference on Pattern*
1088 *Languages of Programs*, pages 1–20, 2018.
- 1089 18 Eva AM van Dis, Johan Bollen, Willem Zuidema, Robert van Rooij, and Claudi L Bockting.
1090 Chatgpt: five priorities for research. *Nature*, 614(7947):224–226, 2023.
- 1091 19 Laria Reynolds and Kyle McDonell. Prompt programming for large language models: Beyond
1092 the few-shot paradigm. *CoRR*, abs/2102.07350, 2021. URL: [https://arxiv.org/abs/2102.](https://arxiv.org/abs/2102.07350)
1093 [07350](https://arxiv.org/abs/2102.07350), arXiv:2102.07350.
- 1094 20 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, Quoc Le, and
1095 Denny Zhou. Chain of thought prompting elicits reasoning in large language models. *CoRR*,
1096 abs/2201.11903, 2022. URL: <https://arxiv.org/abs/2201.11903>, arXiv:2201.11903.
- 1097 21 Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani
1098 Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto,
1099 Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. Emergent abilities of large
1100 language models, 2022. URL: <https://arxiv.org/abs/2206.07682>, doi:10.48550/ARXIV.
1101 2206.07682.
- 1102 22 Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan,
1103 and Jimmy Ba. Large language models are human-level prompt engineers, 2022. URL:
1104 <https://arxiv.org/abs/2211.01910>, doi:10.48550/ARXIV.2211.01910.
- 1105 23 Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh.
1106 Autoprompt: Eliciting knowledge from language models with automatically generated
1107 prompts. *CoRR*, abs/2010.15980, 2020. URL: <https://arxiv.org/abs/2010.15980>, arXiv:
1108 2010.15980.
- 1109 24 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al.
1110 Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- 1111 25 Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale
1112 Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, and Ed Chi. Least-to-most prompting
1113 enables complex reasoning in large language models, 2022. URL: [https://arxiv.org/abs/](https://arxiv.org/abs/2205.10625)
1114 [2205.10625](https://arxiv.org/abs/2205.10625), doi:10.48550/ARXIV.2205.10625.
- 1115 26 Jaehun Jung, Lianhui Qin, Sean Welleck, Faeze Brahman, Chandra Bhagavatula, Ronan Le
1116 Bras, and Yejin Choi. Maieutic prompting: Logically consistent reasoning with recursive
1117 explanations, 2022. URL: <https://arxiv.org/abs/2205.11822>, doi:10.48550/ARXIV.2205.
1118 11822.

- 1119 27 Simran Arora, Avanika Narayan, Mayee F Chen, Laurel Orr, Neel Guha, Kush Bhatia,
1120 Ines Chami, and Christopher Re. Ask me anything: A simple strategy for prompting
1121 language models. In *International Conference on Learning Representations*, 2023. URL:
1122 <https://openreview.net/forum?id=bhUPJnS2g0X>.
- 1123 28 Vivian Liu and Lydia B Chilton. Design guidelines for prompt engineering text-to-image
1124 generative models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing*
1125 *Systems*, pages 1–23, 2022.
- 1126 29 Paula Maddigan and Teo Susnjak. Chat2vis: Generating data visualisations via natural lan-
1127 guage using chatgpt, codex and gpt-3 large language models. *arXiv preprint arXiv:2302.02094*,
1128 2023.
- 1129 30 Xu Han, Weilin Zhao, Ning Ding, Zhiyuan Liu, and Maosong Sun. Ptr: Prompt tuning with
1130 rules for text classification. *AI Open*, 3:182–192, 2022.
- 1131 31 Shuai Wang, Harrisen Scells, Bevan Koopman, and Guido Zuccon. Can chatgpt write a good
1132 boolean query for systematic review literature search? *arXiv preprint arXiv:2302.03495*, 2023.
- 1133 32 Chunqiu Steven Xia and Lingming Zhang. Conversational automated program repair. *arXiv*
1134 *preprint arXiv:2301.13246*, 2023.
- 1135 33 Jonathan H Choi, Kristin E Hickman, Amy Monahan, and Daniel Schwarcz. Chatgpt goes to
1136 law school. *Available at SSRN*, 2023.
- 1137 34 Simon Frieder, Luca Pinchetti, Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiewicz,
1138 Philipp Christian Petersen, Alexis Chevalier, and Julius Berner. Mathematical capabilities of
1139 chatgpt. *arXiv preprint arXiv:2301.13867*, 2023.