

## A Training Examples

Figure 10 shows six training examples from various sources.

## B Anticorrelation between Perplexity and Generation Quality

When fine-tuning LIMA, we observe that perplexity on held-out Stack Exchange data (2,000 examples) negatively correlates with the model’s ability to produce quality responses. To quantify this manual observation, we evaluate model generations using ChatGPT, following the methodology described in Section 5. Figure 9 shows that as perplexity rises with more training steps – which is typically a negative sign that the model is overfitting – so does the quality of generations increase. Lacking an intrinsic evaluation method, we thus resort to manual checkpoint selection using a small 50-example validation set.

## C Human Annotation

Figure 11 shows the human annotation interface we used to collect preference judgments. Annotators were asked to exercise empathy and imagine that they were the original prompters.

To remove the any confounders from the annotation interface, we present anonymized and randomly ordered LIMA response and baseline model response to annotators. Besides, we performed screening annotation tasks with annotators and selected annotators who achieve the highest agreement scores for the final annotation.

## D ChatGPT Score

Automatically evaluating generative models is a difficult problem. For ablation experiments (Section 5), we use ChatGPT (GPT-3.5 Turbo) to evaluate model outputs on a 6-point Likert score given the prompt in Figure 12.

## E Generating Text with Complex Structure

In our preliminary experiments, we find that although LIMA can respond to many questions in our development set well, it cannot consistently respond to questions that specify the structures of the answer well, e.g. summarizing an article into bullet points or writing an article consisting of several key elements. Hence, we investigate whether adding a few training examples in this vein can help LIMA generalize to prompts with unseen structural requirements. We added six examples with various formatting constraints, such as generating a product page that includes *Highlights*, *About the Product*, and *How to Use* or generating question-answer pairs based on a given article.

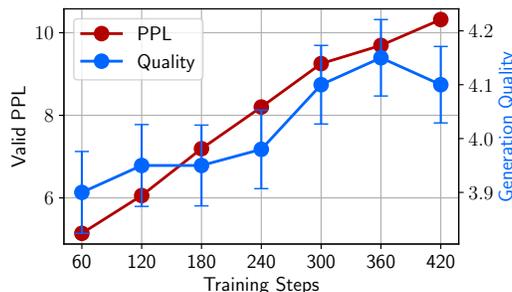


Figure 9: Validation set perplexity versus generation quality (as evaluated by ChatGPT), across the training process of LIMA 65B. We observe similar trends for 7B and 30B parameter models, and across different mixtures of training data.

---

**[Stack Exchange (STEM)]**

**What is the difference between minimum and infimum? I have a great confusion about this.**

The minimum is attained, the infimum isn't necessarily.

Example: Let  $f(x) = \frac{1}{x}$ . Then  $f$  has no minimum value on the interval  $(0, \infty)$ . The minimum is the smallest element in the set. That is  $\min\{f(x) \mid x \in (0, \infty)\}$  doesn't exist because there is not smallest number in the set. Another example is the minimum of the set  $S = (0, 1) = \{x \mid 0 < x < 1\}$ . Here again there isn't a smallest number  $\min\{x \mid 0 < x < 1\}$  doesn't exist.

The infimum of a set  $S$  is defined as the greatest number that is less than or equal to all elements of  $S$  (from Wikipedia). The infimum is also sometimes called the greatest lower bound. It is a fact that every non empty set (bounded below) of real numbers has an infimum. But, as we saw, not every real set has a minimum. So in the example  $\inf\{f(x) \mid x \in (0, \infty)\} = 0$ .

Note that the infimum and the minimum can be the same. Consider for example  $S = \{1, 2, 3, \dots\}$ . Then the infimum and minimum is both 1. Consider this other example. If  $f$  is a continuous function on a closed interval  $[a, b]$ , then it is a fact that  $f$  attains a minimum over that interval. So here again  $\inf\{f(x) \mid x \in [a, b]\} = \min\{f(x) \mid x \in [a, b]\}$ .

---

**[Stack Exchange (Other)]**

**Was the Millennium Falcon a one-off or was it mass produced?**

The Millennium Falcon is not unique; it was a model YT-1300f Corellian light freighter. This implies that there are other similar ships, including other variations on the YT-1300 model. You can find a lot of information about the Millennium Falcon in a companion book, Star Wars: The Force Awakens: Incredible Cross-Sections. The model number also appears on-screen at least once: in Star Wars: Rebels episode "Fighter Flight" on a poster. Note, however, that the Millennium Falcon has been highly customized; some of the custom work is mentioned in The Force Awakens. It's not clear just how much of the ship is stock and how much is custom.

If you're willing to go outside Canon into Legends, the model number of the ship, and other similar Corellian light freighters, can be seen in some of the video games, and there is an entire novel called Millennium Falcon that details its history. The ship was serial #YT-1300 49272ZED, one of many such vessels, but from the beginning the ship was shown to have some "quirks" to distinguish it from its counterparts. In Legends, the ship was also destroyed and rebuilt from scratch at least once, so it's arguable as to whether it's still the same ship or not.

---

**[wikiHow]**

**How to be a lazy college student?**

In college, unlike in high school, you are expected to take responsibility for your own learning. It is ultimately up to you to determine how much time and effort you want to put into your schoolwork. Experts suggest that college students today work considerably fewer hours than their counterparts did in 1961. But there is also evidence that working more hours is not necessarily the best way to be productive. A little laziness may even be good for us.

**Part 1 of 3: prioritize what's important**

1. Ask yourself what you want from college. There are many reasons why a college degree is desirable, not least that it is still considered highly valuable by many employers. Your own reasons for pursuing a college degree may be more personal, but spend some time working out exactly what these reasons are. Consider why college is important to you. Do you want to learn important skills like critical thinking and problem-solving? Do you want to make friends, network and get to know other people? Do you want to get good grades for further study? Do you want to get a particular job that requires a college degree? Spend some time reflecting on these questions and come up with your own answers. It might even be a good idea to write them down. This will help you to recognize what's important to you so that you can focus on the essential things.
2. Select your major and your classes carefully. [...]

---

**[Manually Authored – Chitchat]**

**Tell me an interesting fact about geography.**

Here are some interesting facts about geography:

- Bhutan is one of the most intriguing places in the world. The Buddhist kingdom is the only carbon-negative country on Earth, with strict environmental laws and cultural norms that preserve its stunning natural landscape. For similar reasons, Bhutan was one of the last countries to introduce television in 1999, and has no traffic lights to this day.
- There are currently few territories on Earth that are unclaimed by any sovereign state (terra nullius). One example is Bir Tawil, along the border between Egypt and Sudan, which was created by a discrepancy between borders drawn in 1899 and 1902. The territory remains unclaimed by either country, because it would entail forfeiting the much larger Halaib Triangle, which is claimed by both countries. Bir Tawil has no settled population, but the land is used by Bedouins who roam the area.
- The antipode of any spot on Earth [...]

---

**[Manually Authored – Advice]**

**I'm going to NeurIPS for the first time to present the first paper I've ever published. I'm afraid I'll be alone and overwhelmed by everyone. What should I do?**

Going to an academic conference for the first time may seem like a bewildering experience, but if done right, conferences are a wonderful opportunity to brainstorm with experts from around the world, present your work and get feedback, and make new friends. Here are some tips for improving your conference experience:

- Reach out in advance to people who have inspired your work. [...]
- Show curiosity about other people's work. [...]
- Sign up for student volunteer [...]
- Ask your advisor to introduce you [...]
- At the same time, try to socialize outside of your lab [...]

Ultimately, it is important to remember that every first-conference attendee is probably feeling a bit shy too, and that taking the first step in starting a conversation will almost always be met with a big smile!

---

**[Manually Authored – Writing]**

**I am planning to start a book club with some friends. Can you write an email invitation for the same?**

Here is an example email invitation for a book club:

Subject: Join me for a book club!

Hi [Friend's Name], how are you?

I'm starting a book club with some friends, and I'd love to have you join us!

The book club is a casual and fun way to read and discuss books that we love or want to learn more about, and also make new friends. We plan to meet more-or-less once a month.

We're thinking of having our first meeting on [Date & Time]. Could that work for you?

Would love to see you there!

Thanks,  
[Your Name]

---

Figure 10: **Training examples from different sources.** *Top row:* examples mined from community Q&A. *Bottom row:* manually-authored examples. The blue text is the prompt. The gray ellipsis [...] indicates that the response was trimmed to fit this page, but the actual training example is longer.

After training with these six additional examples, we test the model on a few questions with format constraints and observe that LIMA responses greatly improve. We present two examples in Figure 13, from which we can see that LIMA fails to generate proper answers without structure-oriented training examples (left column), but it can generate remarkably complex responses such as a marketing plan even though we do not have any marketing plan examples in our data (right column).

---

Imagine that you have a super-intelligent AI assistant, and that you require help with the following question. Which answer best satisfies your needs?

**Question:** <QUESTION>

**Answer A:**

<ANSWER A>

**Answer B:**

<ANSWER B>

**Comparing these two answers, which answer is better?**

- Answer A is significantly better.
- Answer B is significantly better.
- Neither is significantly better.

---

Figure 11: Human annotation interface.

---

You are evaluating a response that has been submitted for a particular task, using a specific set of standards. Below is the data:

[BEGIN DATA]

\*\*\*

[Task]: {task}

\*\*\*

[Submission]: {submission}

\*\*\*

[Criterion]: helpfulness:

"1": "Not helpful - The generated text is completely irrelevant, unclear, or incomplete. It does not provide any useful information to the user."

"2": "Somewhat helpful - The generated text has some relevance to the user's question, but it may be unclear or incomplete. It provides only partial information, or the information provided may not be useful for the user's needs."

"3": "Moderately helpful - The generated text is relevant to the user's question, and it provides a clear and complete answer. However, it may lack detail or explanation that would be helpful for the user."

"4": "Helpful - The generated text is quite relevant to the user's question, and it provides a clear, complete, and detailed answer. It offers additional information or explanations that are useful for the user. However, some of the points of the response are somewhat repetitive or could be combined for greater clarity and concision"

"5": "Very helpful - The generated text is highly relevant to the user's question, and it provides a clear, complete, and detailed answer. It offers additional information, explanations, or analogies that are not only useful but also insightful and valuable to the user. However, the structured of the response is not well-organized and there is no clear progression or logical sequence of different points in the response."

"6": "Highly helpful - The generated text provides a clear, complete, and detailed answer. It offers additional information or explanations that are not only useful but also insightful and valuable to the user. The response is also in a logical and easy-to-follow manner by explicitly using headings, bullet points, or numbered lists to break up the information and make it easier to read."

\*\*\*

[END DATA]

Does the submission meet the criterion? First, write out in a step by step manner your reasoning about the criterion to be sure that your conclusion is correct. Avoid simply stating the correct answers at the outset. Then print the choice only from "1, 2, 3, 4, 5, 6" (without quotes or punctuation) on its own line corresponding to the correct answer. At the end, repeat just the selected choice again by itself on a new line.

---

Figure 12: Prompt for ChatGPT evaluation with a 6-scale Likert score. The placeholders "task" and "submission" will be replaced by specific details from the actual case being evaluated.

LIMA w/o 6 Format Constraint Examples (994 Examples)	LIMA (1,000 Examples)
<p><b>[Marketing Plan (Out of Distribution)]</b>  <b>I want to start a local coffee shop. Create a marketing plan with the following elements: Marketing Goals and Objectives, Define Target Audience, Research Marketing Tactics, Plan Marketing Tactics and Develop Your Timeline and Budget.</b></p>	
<p>Coffee Shop Marketing Plan</p> <p>Executive Summary</p> <p>The following marketing plan for Java Jolt Coffee Bar outlines an aggressive approach to establishing a thriving coffee shop business in downtown Pittsburgh. Java Jolt will capitalize on the popularity of specialty coffees and the growing demand for WiFi hot spots by offering a comfortable environment where customers can enjoy premium coffee drinks while taking advantage of free Internet access. Market research indicates that there is a significant opportunity for a new coffee bar located in the heart of the downtown business district. Market trends suggest that consumers are willing to pay a higher price for quality specialty coffees served in a welcoming atmosphere. Additionally, many office workers and students living in the area do not have easy access to high-speed Internet service, presenting an additional revenue stream for the coffee bar.</p> <p>This marketing plan includes the following elements:</p> <ul style="list-style-type: none"> <li>• Market Research supporting the need for a new coffee shop in the target location</li> <li>• Detailed Description of the products and services offered by Java Jolt</li> <li>• Target Market Analysis identifying the primary customer segments for Java Jolt</li> <li>• Marketing Strategy including pricing decisions, promotional activities, and ways to differentiate Java Jolt from competitors</li> <li>• Implementation Plan detailing the actions necessary to launch Java Jolt successfully</li> <li>• Budget allocating funds to each component of the marketing plan</li> </ul> <p>A comprehensive marketing plan has been developed based on extensive research into the local market conditions, consumer trends, and competitive landscape. With its unique combination of premium coffee drinks and wireless Internet access, Java Jolt will quickly become a favorite destination for people seeking a relaxing environment to socialize, study, or conduct business.</p>	<p>Here's a marketing plan for your local coffee shop.</p> <p><b>Marketing Goals and Objectives</b></p> <ul style="list-style-type: none"> <li>• Increase awareness of our brand among consumers [...]</li> <li>• Drive foot traffic [...] during slow business hours [...]</li> <li>• Build relationships with influencers in our community who [...]</li> </ul> <p><b>Define Target Audience</b></p> <p>Our target audiences include:</p> <ul style="list-style-type: none"> <li>• Residents of the neighborhood surrounding our coffee shop</li> <li>• Office workers near our coffee shop</li> <li>• Students at nearby schools</li> </ul> <p>We will focus our marketing efforts on these audiences because they are most likely to become regular customers at our [...]</p> <p><b>Research Marketing Tactics</b></p> <p>Here are some marketing tactics that we could use to reach to reach our target audiences and achieve our goals:</p> <ul style="list-style-type: none"> <li>• Host events at our coffee shop that appeal to our target audiences (e.g., open mic nights for students or networking events [...])</li> <li>• Offer discounts during slow business hours to entice [...]</li> <li>• [...]</li> </ul> <p><b>Plan Marketing Tactics</b></p> <p>Here's how we will implement the marketing tactics above:</p> <ul style="list-style-type: none"> <li>• Events: We will host one event per month at our coffee shop. For each event, we will send out email invitations [...]</li> <li>• Discounts: Every Monday through Thursday from 3pm–6pm, we will offer \$1 off any drink purchase. To promote [...]</li> <li>• Local partnerships: We will reach out to five local businesses and schools to ask if we can leave stacks of coupons [...]</li> <li>• Social media: [...]</li> <li>• Reviews: [...]</li> </ul> <p><b>Develop Your Timeline and Budget</b></p> <p>This is what our marketing timeline looks like:</p> <ul style="list-style-type: none"> <li>• Month 1: Launch email newsletter and social media accounts</li> <li>• Month 2: Start distributing coupons at [...]</li> <li>• [...]</li> </ul> <p>Here's the budget for each marketing activity:</p> <ul style="list-style-type: none"> <li>• Email newsletter: \$50/month for MailChimp subscription</li> <li>• Social media: \$100/month for paid ads [...]</li> <li>• [...]</li> </ul>
<p><b>[Summarization into Bullet Points (In Distribution)]</b>  <b>Can you summarize this message into bullet points? Gradually slowing job gains and a growing labor force in March [...] Mr. Biden said [...]. More Americans joined the labor force, [...] this spring that he will seek re-election.</b></p>	
<p>Though the unemployment rate fell to its lowest level since May 1970, at 4.8 percent, Mr. Biden still faces an economy that leaves many Americans feeling left behind. Many working-class [...]</p> <p>Mr. Biden's challenge will be persuading those voters that the economic picture has [...] His opponents have already begun painting a bleak portrait of the economy, arguing that most Americans have not benefited from the president's policies.</p> <p>"The March jobs report shows continued solid job growth but also indicates that we must do better," said Representative Nancy Pelosi of California, the Democratic leader in the House. [...]</p>	<p>Here are the main points from the text:</p> <ul style="list-style-type: none"> <li>• Job creation slowed to 236,000 for the month of March, which is close to the level needed to stabilize the economy and prices.</li> <li>• More Americans joined the labor force, and wage gains fell slightly, both of which should help cool inflation.</li> <li>• The report underscores the political tensions facing President Biden ahead of his expected announcement that he will seek re-election.</li> </ul>

Figure 13: Model outputs from test prompts that ask the model to generate according to specified structures. The gray ellipsis [...] indicates that the response was trimmed to fit this page, but the generated text is actually longer.