

L3Ms — LAGRANGE LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Supervised fine-tuning (SFT) and alignment of large language models (LLMs) are key steps in providing a good user experience. However, the concept of an appropriate alignment is inherently application-dependent, and current methods often rely on heuristic choices to drive the optimization. In this work, we formulate SFT and alignment as a constrained optimization problem: the LLM is fine-tuned on a task while being required to meet application-specific requirements, without resorting to heuristics. To solve this, we propose Lagrange Large Language Models (L3Ms), which employ logarithmic barriers to enforce the constraints. This approach allows for the customization of L3Ms across diverse applications while avoiding heuristic-driven processes. We demonstrate experimentally the versatility and efficacy of L3Ms in achieving tailored alignments for various applications.

1 INTRODUCTION

Large language models (LLMs) are used for a wide range of tasks: as chatbots (Brown et al., 2020; OpenAI, 2024), for code generation (Ahmad et al., 2021; Wang et al., 2021; Rozière et al., 2024), for medical assistance (Yang et al., 2022; Moor et al., 2023), and more. The key ingredients for their impressive downstream performance are supervised fine-tuning (SFT) and alignment; the former fine-tunes the LLM to a task of interest, while the latter instills it with preferential properties. Arguably, the *right* combination of preferential properties is highly application/task-dependent. For instance, a scholar might want a chatbot to be honest and factual to assist with their work, whereas a fiction writer might prefer the opposite behavior to help create fantastical imaginary worlds. There is also plenty of (anecdotal) evidence in support: some LLMs refuse to provide information on how to “kill” a process in Unix, recommending the use of less violent strategies for dealing with wayward computer programs instead.¹ Therefore, we need frameworks for the customization of LLMs.

Consequently, Li et al. (2021); Bai et al. (2022); Rame et al. (2023); Wu et al. (2023); Ji et al. (2023) fine-tune LLMs on varying combinations of such preferential properties. In practice, one tends to resort to trial and error to find the right preference combination for their particular task. In doing so, one verifies if certain *minimum baselines* are satisfied, such as ensuring the factual correctness of statements or confirming that the response lengths are capped at 100 words. Since there isn’t a way to enforce such requirements directly, current methods resort to heuristics. Additionally, existing pipelines carry out SFT and alignment in sequential steps. Hence they need to ensure that the LLM does not forget relevant task information learned during the SFT stage. This is achieved by penalizing the LLM for drastic deviations: the strength of the penalty is determined heuristically.

In this work, we formulate SFT and alignment in LLMs as a constrained optimization problem. In particular, we fine-tune an LLM to minimize the task objective (the objective function) while simultaneously satisfying application-specific minimum requirements (the constraints). This merges the SFT and alignment stages and mitigates the reliance on heuristics altogether. Furthermore, we propose Lagrange Large Language Models, a.k.a. L3Ms, to solve such constrained optimization problems. Specifically, they do so by employing logarithmic barriers and enforcing the constraints gradually over the optimization procedure. Lastly, we empirically demonstrate how one can pick and choose constraints and tailor L3Ms to a range of applications without resorting to heuristics. For example, Table 1 provides the generated responses from two such L3Ms; both fine-tuned on the task of instruction-following, but one is constrained to be concise while the other is to be verbose.

¹https://www.reddit.com/r/LocalLLaMA/comments/180p17f/new_claude_21_refuses_to_kill_a_python_process/

054	Prompt	What are the four main types of macromolecules found in living organisms?
055		
056	Response	1. Nucleic acids (DNA and RNA)
057	(short response L3M)	2. Proteins
058		3. Carbohydrates
059		4. Lipids
060		
061	Response	1. Nucleic acids, including DNA and RNA, which carry genetic information
062	(long response L3M)	2. Proteins, which are responsible for the structure, function, and regulation of living organisms
063		3. Lipids, which provide energy and form the membranes that surround cells
064		4. Carbohydrates, which are used as an energy source and as structural components of cells.
065		
066		
067	Prompt	Who discovered the law of universal gravitation?
068		
069	Response	Sir Isaac Newton is credited with discovering the law of universal gravitation in the late 17th century.
070	(short response L3M)	
071		
072	Response	Sir Isaac Newton is widely credited with the discovery of the law of universal gravitation, which states that every particle in the universe attracts every other particle with a force proportional to their masses and inversely proportional to the square of the distance between them.
073	(long response L3M)	
074		

Table 1: **Example responses from length constrained L3Ms.** We provide example responses from L3Ms with varying length constraints. We include the prompt along with the generated responses from two L3Ms: one constrained to have short responses, and the other constrained to long ones.

In summary, our contributions are as follows:

1. We formulate SFT and alignment in LLMs as a constrained optimization problem: an LLM is fine-tuned on a task while simultaneously satisfying custom requirements (cf. Section 4).
2. We propose L3Ms, a family of LLMs fine-tuned using the above framework (cf. Section 5).
3. We demonstrate experimentally how L3Ms can be customized to different applications and their specific requirements while avoiding heuristic-driven processes (cf. Section 6).

2 OPTIMIZATION FOR LLMs

Training an LLM proceeds in multiple stages (Ouyang et al., 2022), which we discuss below.

2.1 PRE-TRAINING

The pre-training stage instills the LLM with a generic knowledge of language. It entails regenerating text/token sequences by minimizing their perplexity, i.e., the negative log-likelihood of the sequence normalized by its length. More formally, the perplexity on a sequence x is defined as:

$$l_{\theta}(x) = -\frac{\log \pi_{\theta}(x)}{|x|} = -\frac{1}{|x|} \sum_{i=1}^{|x|} \log \pi_{\theta}(x_i | x_{<i}),$$

where x_i and $x_{<i}$ denote the i -th token in the sequence and its prefix, respectively. Additionally, the function $\pi_{\theta}(\cdot)$ denotes the LLM’s predicted probability distribution over token sequences, where the LLM is parameterized with weights θ . Then, the pre-training objective is given as follows:

$$\min_{\theta} \mathbb{E}_{x \sim q(\cdot)} [l_{\theta}(x)],$$

and the expectation is replaced by an empirical average over a large corpus with trillions of tokens.

2.2 SUPERVISED FINE-TUNING (SFT)

Next, one fine-tunes the LLM to a task of interest, such as instruction-following, summarization, or translation. The data are (prompt, response) pairs (x, y) , and the LLM is fine-tuned to regenerate the responses. Thus, one minimizes the perplexity (or a related loss) on the response given the prompt:

$$\min_{\theta} \mathbb{E}_{(x,y) \sim p(\cdot)} [l_{\theta}(y|x)], \quad (1)$$

for the distribution $p(\cdot)$ over (prompt, response) pairs which reflects the task-related data.

2.3 ALIGNMENT

This stage aligns the LLM to generate responses with preferential properties. A common setup is to learn preference reward functions that represent properties like helpfulness and harmlessness (Bai et al., 2022), followed by reinforcement learning to adapt the LLM to maximize said rewards. This is referred to as reinforcement learning from human feedback (RLHF; Knox & Stone, 2008; Griffith et al., 2013; Christiano et al., 2017). Note that the preference reward functions need not always be learned; they could also be engineered or rule-based, such as the length of the response.

Given a single preference reward function $r(\cdot)$, the alignment objective is given as follows:

$$\max_{\theta} \mathbb{E}_{\substack{(x,\cdot) \sim p(\cdot) \\ y \sim \pi_{\theta}(\cdot|x)}} [r(y|x)]. \quad (2)$$

This maximizes the rewards for the LLM’s responses to prompts sampled from the task distribution. To prevent over-optimization of the reward and avoid drastic deviation away from the SFT model, it is common practice to add a regularization penalty such as the KL divergence (Gao et al., 2023).

We are interested in $k \geq 1$ different preferential properties. In such a scenario, one could learn individual preference reward functions r_i ’s and optimize the LLM to maximize their combination. In particular, Li et al. (2021); Rame et al. (2023); Wu et al. (2023); Ji et al. (2023) use a linear combination of the rewards, substituting the single reward in Eq. (2) with $\sum_{i=1}^k \alpha_i r_i(y|x)$, for some choice of $\alpha_i \geq 0$. Alternatively, one could learn a single reward function by choosing the data proportions of the different properties used to train it. For instance, Bai et al. (2022) use a 3:1 proportion of helpfulness to harmlessness data to train a single preference reward function. Note that the proportions can also be represented via weights $\alpha_i \geq 0$ with $\sum_{i=1}^k \alpha_i = 1$. Therefore, the choice of α_i ’s, combined with the strength of the regularization penalty, together steer the alignment. They are chosen based on the good judgment of the practitioner in a somewhat heuristic manner.

2.4 SHORTCOMINGS

While the above pipeline is commonly used to train LLMs for deployment, it has shortcomings.

Firstly, how does one choose the weights α_i ’s? Or equivalently, *what is the right combination of preference properties?* Choosing between properties such as truthfulness, verbosity, and humor depends on the application at hand. However, even when the application is known, the weights are chosen through trial and error: trying different combinations and inspecting the achieved rewards. For instance, if one wants the responses to be under 100 words for a summarization task, one might repeatedly examine the length of the responses and adjust the weights α_i ’s to achieve that. This inherently involves verification against a set of minimum baselines being satisfied. *Can we avoid heuristics and enforce such minimum requirements for the preference rewards in a principled way?*

Secondly, recall that the original task objective is in Eq. (1). However, optimizing a special purpose objective like Eq. (2) can lead to a decrease in performance on the original task. While penalizing the deviation away from the SFT model helps mitigate this to some extent, the strength of the penalty is again determined heuristically. *Can we ensure the performance on the original task is maintained?*

3 RELATED WORK

3.1 ONE SIZE DOES NOT FIT ALL

Early work on LLM alignment assumed homogeneity of preferences (Bakker et al., 2022). However, the reality is quite different: human preferences vary widely and are highly application-specific

(Rame et al., 2023; Casper et al., 2023). Consequently, Rame et al. (2023); Wu et al. (2023); Ji et al. (2023); Zhou et al. (2024) linearly combine several fine-grained preference rewards, where each combination reflects a different customized preference. However, the *right* combination of preferential properties for a particular application is determined heuristically through trial and error. Conversely, we formulate alignment as a constrained optimization problem, where the custom requirements are naturally imposed as constraints. This alleviates the need for heuristic choices.

3.2 CONSTRAINED ALIGNMENT FOR LLMs

Independent of our work, Moskovitz et al. (2024); Dai et al. (2024) also introduced constrained optimization for LLM alignment but for varying reasons. Motivated to avoid over-optimization of preference rewards, Moskovitz et al. (2024) find “proxy points”, values of the reward functions beyond which the performance of the LLM is negatively impacted. They constrain the average rewards to be in the vicinity of these proxy points. Dai et al. (2024) trade-off between helpfulness and harmlessness from a safety standpoint. They maximize the LLM’s helpfulness while constraining its average harmlessness reward. Our motivation on the other hand is to tailor to custom preferences through sets of different constraints, while simultaneously learning to perform the task at hand.

Our work is different in two important ways. Firstly, both Moskovitz et al. (2024); Dai et al. (2024) consider the alignment process in isolation, because of which their objective is either to maximize one of the preference rewards or to minimize the deviation away from the SFT model. On the other hand, we merge the SFT and alignment stages and choose to minimize the task objective as our objective. This directly ensures that the LLM’s task-solving capabilities are preserved, and it avoids loading an additional reference (SFT) model during the fine-tuning process to compute deviations.

Secondly, both Moskovitz et al. (2024); Dai et al. (2024) obtain solutions to the constrained problem by computing the saddle-point of the Lagrangian. This is done by formulating a min-max game where the LLM minimizes the Lagrangian and the Lagrange multipliers adapt to maximize it. However, the non-convexity of the objective makes this nontrivial. Instead, we choose the logarithmic barrier approach as the associated Lagrange multipliers satisfy the KKT complementary slackness condition by design (cf. Section 5.2) rather than letting the learning procedure do so (which in our experience is extremely sensitive to the choice of the learning rate). We empirically observe that while both methods satisfy the constraints, our approach using log-barriers minimizes the task objective better (cf. Section 6.3). As a side effect, we avoid introducing new learning parameters.

Additionally, we provide experimental validation of our approach on more preferences than Moskovitz et al. (2024); Dai et al. (2024), and use a $4.5\times$ larger LLM than Moskovitz et al. (2024).

4 CONSTRAINED OPTIMIZATION FOR LLMs

Our goal is to reduce the task objective in Eq. (1) for the LLM to solve the intended task, while also enabling custom alignment by having the LLM meet the application-specific minimum requirements on different preference rewards. To do so, we propose the constrained optimization problem:

$$\min_{\theta} \mathbb{E}_{(x,y)\sim p(\cdot)} [l_{\theta}(y|x)] \quad (3)$$

$$\text{subject to } \mathbb{E}_{\substack{(x,\cdot)\sim p(\cdot) \\ y\sim\pi_{\theta}(\cdot|x)}} [r_i(y|x)] \geq b_i \text{ for all } i \in \{1, 2, \dots, k\}. \quad (4)$$

Here, the objective is the same as the task objective in SFT, and the constraints are enforced to satisfy the custom requirements; this merges the SFT and alignment stages. The b_i ’s signify the minimum baselines for each preference reward function $r_i(\cdot)$ and the constraints are enforced on average.

Compared to the previous approach, we no longer rely on heuristics to find a set of weights α_i ’s to satisfy the minimum requirements; we can do so directly through the constraints. Furthermore, with the same objective as SFT, we can directly maintain task performance without any deviation penalty. Additionally, note that whenever a constraint is satisfied, its influence vanishes. For instance, if the LLM is naturally harmless and $r_{\text{harmless}}(y|x) \geq b_{\text{harmless}}$, then the constraint is not active and the LLM will not be penalized. Conversely, the previous approach would further penalize the LLM.

Notation For ease of notation, we rewrite the constrained problem in Eqs. (3) and (4) as follows:

$$\min_{\theta} L(\theta) \text{ subject to } C_i(\theta) \leq 0 \text{ for all } i \in \{1, 2, \dots, k\}, \quad (5)$$

with the objective $L(\theta) = \mathbb{E}_{(x,y) \sim p(\cdot)}[l_{\theta}(y|x)]$ and constraints $C_i(\theta) = \mathbb{E}_{\substack{(x,\cdot) \sim p(\cdot) \\ y \sim \pi_{\theta}(\cdot|x)}}[b_i - r_i(y|x)]$.

4.1 TYPES OF CONSTRAINTS

While we write the constraints in Eq. (4) as expectation/average constraints, other forms exist. For instance, uniform constraints impose a minimum reward on every generated (prompt, response) pair:

$$r_i(y|x) \geq b_i \text{ for all } (x, \cdot) \sim p(\cdot), y \sim \pi_{\theta}(\cdot|x) \text{ and all } i \in \{1, 2, \dots, k\}. \quad (6)$$

Additionally, chance constraints bound the probability of the inequality holding away from zero:

$$\mathbb{P}_{\substack{(x,\cdot) \sim p(\cdot) \\ y \sim \pi_{\theta}(\cdot|x)}}[r_i(y|x) \geq b_i] \geq 1 - \epsilon_i \text{ for all } i \in \{1, 2, \dots, k\}. \quad (7)$$

While these constraints are not equivalent, they are related. We can rewrite Eq. (7) in the form of expectation constraints by using $1 - \epsilon_i$ as the threshold and taking the expectation of the indicator $\mathbb{1}\{r_i(y|x) \geq b_i\}$. Moreover, Eq. (6) implies Eq. (4), but the converse is not true. Unfortunately, Eq. (6) is difficult to achieve in practice, especially when the data distribution is unknown.

We continue using the expectation constraints, but similar discussions can extend to the other types.

4.2 LAGRANGE MULTIPLIERS

We can introduce Lagrange multipliers $\lambda_i \geq 0$ for the constraints and obtain the Lagrangian:

$$\mathcal{L}(\theta) = L(\theta) + \sum_{i=1}^k \lambda_i C_i(\theta). \quad (8)$$

There is a rich literature connecting the Lagrangian with constrained optimization. Notably, the KKT conditions (Karush, 1939; Kuhn & Tucker, 1951) provide sufficiency conditions for global optimality under convexity, where the solution can be obtained by finding the saddle point of the Lagrangian. However, these conditions are not enough for highly non-convex scenarios like ours.

Nonetheless, the Lagrangian is instructive in understanding the relative importance of the constraints. For an active constraint, i.e., one satisfied with equality, the corresponding Lagrange multiplier can be non-zero; the larger its value, the more important the constraint. Conversely, for an inactive constraint, i.e., one satisfied with strict inequality, the corresponding Lagrange multiplier must vanish to 0. This is known as complementary slackness and is one of the KKT conditions.

4.3 LOGARITHMIC BARRIER

A practical way to enforce constraints is with barrier functions. Consider the (relaxed) log-barrier:

$$\mathcal{B}_{\mu,s}(z) = \begin{cases} -\mu \log(-z), & z \leq -s \\ \frac{\mu}{s}z + \mu - \mu \log s, & z > -s \end{cases}, \text{ and hence } \partial_z \mathcal{B}_{\mu,s}(z) = \frac{\mu}{\max(-z, s)}, \quad (9)$$

with parameters $\mu, s > 0$. This is a convex, continuous, and differentiable function, which is valid for all $z \in \mathbb{R}$. Importantly, for $s = \mu^2$, this barrier function converges to the characteristic function $\chi\{z \leq 0\}$ as $\mu \rightarrow 0$, i.e., it takes the value 0 when $z \leq 0$ and ∞ otherwise (Tal et al., 1992; Nash et al., 1994; Hauser & Saccon, 2006; Feller & Ebenbauer, 2017); the condition $s = \mu^2$ is sufficient, but not necessary (Kervadec et al., 2022). This convergence to the characteristic function is visually depicted in Fig. 1, showing the change in the log-barrier function as we gradually decrease μ .

We can now use the log-barrier to enforce the constraints in Eq. (5) and simply add them to the objective. We obtain an *unconstrained* objective, with μ controlling the strength of the constraints:

$$\mathcal{G}_{\mu}(\theta) = L(\theta) + \frac{1}{k} \sum_{i=1}^k \mathcal{B}_{\mu,\mu^2}(C_i(\theta)). \quad (10)$$




Figure 1: **The relaxed logarithmic barrier.** We depict the convergence of the relaxed logarithmic barrier $\mathcal{B}_{\mu, \mu^2}(z)$ to the characteristic function $\chi\{z \leq 0\}$ as $\mu \rightarrow 0$. We gradually decrease μ from 1 (blue) to 0.01 (red). Consequently, $\mathcal{B}_{\mu, \mu^2}(z)$ gets closer to 0 for $z \leq 0$ and increases to ∞ otherwise.

5 LAGRANGE LARGE LANGUAGE MODELS (L3MS)

We have thus far formulated the SFT and alignment stages as a constrained optimization problem in Eq. (5). We proceed to find solutions for the same by solving the unconstrained objective in Eq. (10). We call the family of models obtained in this way L3Ms, i.e., Lagrange Large Language Models.

5.1 OPTIMIZATION PROCEDURE

Since the log-barrier converges to the characteristic function as $\mu \rightarrow 0$, we want to find the minimizer of $\mathcal{G}_\mu(\theta)$ for a very small μ . However, doing so directly leads to instabilities as the objective function is ill-conditioned. Instead, it is common practice to follow an iterative procedure where one finds the minimizer for a fixed μ , reduces μ , and repeats (Curtis et al., 2024). Specifically, the procedure is instantiated with initial values θ_0 , μ_1 , and $0 < \gamma < 1$. On the t -th iteration, $\mu_t \leftarrow \gamma \mu_{t-1}$ is reduced and $\theta_t \leftarrow \arg_\theta \min \mathcal{G}_{\mu_t}(\theta)$ (with initialization at θ_{t-1}). In doing so, the constraints are enforced gradually, nudging the LLM to satisfy them over the optimization procedure while avoiding instabilities. As $\{\mu_t\} \searrow 0$, the weights θ_t converge to the minimizer of the constrained problem.

It is impossible to minimize $\mathcal{G}_{\mu_t}(\theta)$ exactly in many practical applications. Instead, at each iteration, we take a single optimization step toward the solution. Doing so is amenable to stochastic gradient methods and mitigates computational overhead — the optimization proceeds as normal while the value of μ is reduced over the course of the procedure. One can guarantee convergence of such a procedure to the optimal solution in some settings; for instance, Curtis et al. (2024) prove convergence when dealing with box constraints. However, convergence in a scenario like ours is not guaranteed. Nonetheless, we will experimentally demonstrate its use for our constrained problems.

We employ stochastic gradient methods and derive the gradient of our objective function directly:

$$\partial_\theta \mathcal{G}_\mu(\theta) = \partial_\theta L(\theta) + \frac{\mu}{k} \sum_{i=1}^k \frac{\partial_\theta C_i(\theta)}{\max(-C_i(\theta), \mu^2)}. \quad (11)$$

This follows immediately from Eqs. (9) and (10). Note that the $\partial_\theta C_i(\theta)$'s are also known as the policy gradients in reinforcement learning literature. We discuss our strategy for estimating these gradients in Appendix B and refer readers to Schulman et al. (2016) for a more detailed review.

5.2 CONNECTION TO LAGRANGE MULTIPLIERS

The log-barrier and the Lagrangian are intrinsically connected; this becomes evident when comparing Eq. (11) with the (gradient of the) Lagrangian in Eq. (8). In particular, we define the multipliers:

$$\hat{\lambda}_i = \frac{\mu}{k \max(-C_i(\theta), \mu^2)},$$

324 corresponding to the gradients $\partial_{\theta} C_i(\theta)$'s in Eq. (11). They can be interpreted as Lagrange multi-
 325 pliers: for active constraints, $\hat{\lambda}_i = 1/k\mu$ is non-zero; for inactive constraints, $\hat{\lambda}_i = -\mu/kC_i(\theta)$
 326 vanishes to 0 as $\mu \rightarrow 0$. Hence, the KKT complementary slackness condition is satisfied by design.
 327

328 5.3 IMPLEMENTATION DETAILS

329
 330 **Alternating objective and gradient clipping** A simple, yet effective way to ensure stable training
 331 of large models is gradient clipping (Goodfellow et al., 2016; Zhang et al., 2020). We employ this
 332 technique, albeit with the modification of clipping the gradients of both the task objective and the
 333 constraints separately, as they can have varying magnitudes. We achieve this by alternating between
 334 reducing the task objective and enforcing the constraints by flipping a fair coin to select one or the
 335 other. While this doubles the number of steps to achieve the same effect, it does not increase the
 336 amount of work done as now only one part of the objective or the other is evaluated at each step.
 337

338 **Length normalization** The gradient of our objective function in Eq. (11) involves the log-
 339 likelihoods of the generated responses through the gradients $\partial_{\theta} C_i(\theta)$'s (cf. Eq. (12)). To avoid
 340 a response length bias, we length-normalize the log-likelihoods, akin to the definition of perplexity.
 341

342 **Estimating the mean preference rewards** We need to estimate the expectations involved in the
 343 gradient of our objective function in Eq. (11). The expectations in the numerators can be estimated
 344 with the per-mini-batch Monte Carlo averages (Mohamed et al., 2020). However, $C_i(\theta)$ in the de-
 345 nominator needs careful consideration. Note that: (i) $C_i(\theta)$ does not involve the gradient, so its
 346 estimates can include information from previous mini-batches to reduce variance, and (ii) since the
 347 weights are updated during fine-tuning, $C_i(\theta)$ is non-stationary. Consequently, we use an exponen-
 348 tial moving average estimate for the mean (offset) preference rewards $C_i(\theta)$ in the denominator.
 349

350 6 EXPERIMENTAL RESULTS

351
 352 To illustrate the versatility of L3Ms, we empirically evaluate them on two criteria: (i) the satisfaction
 353 of the imposed custom constraints, and (ii) the minimization of the task objective. Note that L3Ms
 354 have essentially the same time complexity as traditional approaches: the SFT and alignment stages
 355 have been combined and the log-barrier parameter is adjusted during the fine-tuning procedure itself.
 356

357 6.1 SETUP

358
 359 We use LLaMA-7B (Touvron et al., 2023) for all our experiments, as it is a lightweight LLM pre-
 360 trained on a large corpus. We are interested in the instruction-following task, for which we use
 361 UltraChat (Ding et al., 2023), a large-scale dataset of instructional conversations. We refer to the
 362 model fine-tuned via the SFT objective (without alignment) as the *SFT* model. We fine-tune LLMs
 363 using the min-max approach to find the saddle-point of the Lagrangian, as proposed by Moskovitz
 364 et al. (2024); Dai et al. (2024), referring to them as *MMs*. Lastly, we refer to models fine-tuned using
 365 our approach as *L3Ms*. We use the Transformers package (Wolf et al., 2020) for our implementation
 366 and run experiments on NVIDIA H100s. Further details on our setup are included in Appendix A.

367 In what follows, we use different preference reward functions and vary the custom constraint require-
 368 ments. All results are obtained on a held-out test dataset (not seen during training or validation).
 369

370 6.2 LENGTH CONSTRAINED L3MS

371
 372 Consider tasks where the lengths of the responses need to be contained in the range $[l_{\text{low}}, l_{\text{high}}]$ to
 373 control its verbosity, in summarization tasks for example (Makino et al., 2019). The natural choice
 374 of reward functions in this case are the ones that compute the response length and its negation via
 375 $r_1(y|x) = |y|$ and $r_2(y|x) = -|y|$; correspondingly, these rewards are to be controlled with the
 376 requirements of l_{low} and l_{high} respectively. Note that these rewards are perfectly anti-correlated.
 377

If we naively average the rewards, any unconstrained formulation of alignment (including RLHF)
 will be ineffective as the loss will always vanish due to the anti-correlation. We could use a weighted

LLM type	Length	Perplexity
SFT	121.6	0.805 \pm 0.3
L3M [50 , 100]	81.3	0.804 \pm 0.3
L3M [100 , 150]	120.7	0.804 \pm 0.3
L3M [50 , 75]	64.4	0.807 \pm 0.3
L3M [75 , 100]	88.2	0.808 \pm 0.3
L3M [100 , 125]	111.7	0.810 \pm 0.3
L3M [125 , 150]	126.5	0.809 \pm 0.3
L3M [75 , 87.5]	82.9	0.811 \pm 0.3
L3M [87.5, 100]	92.7	0.809 \pm 0.3
L3M [100 , 112.5]	104.8	0.810 \pm 0.3
L3M [112.5, 125]	117.3	0.810 \pm 0.3




Figure 2: **Length constrained L3Ms.** We report the response lengths (in tokens) and task perplexities of the SFT model and the L3Ms with varying length constraints. *Left:* The mean response length with the mean and standard deviation of the task perplexities. *Right:* The distribution of the response lengths. The notches indicate the medians and their 95% confidence intervals, the boxes show the $\pm 25\%$ quantiles, and the whiskers denote the $1.5\times$ interquartile ranges. The white circles mark the means and the black dashed lines depict the constraints imposed on the different L3Ms.

average and tune the weights heuristically, but this is tedious. Instead, we use the constrained formulation and directly constrain the rewards $r_1(y|x) = |y| \geq l_{\text{low}}$ and $r_2(y|x) = -|y| \geq -l_{\text{high}}$.

We fine-tune several L3Ms with varying length constraints. We illustrate the distributions of the generated response lengths (in tokens) and report the perplexities achieved on the task-related data in Fig. 2. We observe that the mean response lengths are in the required range in each case, satisfying the imposed expectation constraints. Additionally, the task perplexities increase slightly as the constraints on the response lengths are made more stringent. However, there is little to no degradation relative to the SFT model, with all mean task perplexities being within 0.02 standard deviations.

The examples included in Table 1 are generated from such L3Ms. While all the responses correctly answer the prompts, their lengths vary corresponding to the length constraints imposed on them.

6.3 HELPFUL AND HARMLESS L3MS

Next, we consider the Helpful and Harmless preferences that have been extensively used in the LLM alignment literature (Ji et al., 2023; Wang et al., 2024; Zhou et al., 2024; Guo et al., 2024). Specifically, we utilize the datasets by Bai et al. (2022) to train two preference reward functions respectively. These learned reward functions are negatively correlated (Bai et al., 2022; Dai et al., 2024). Furthermore, note that the numerical outputs of both these reward functions are interpreted as ratings such that a higher numerical value indicates higher helpfulness/harmlessness and vice versa.

We fine-tune several L3Ms with varying HH constraints. We compare our L3M approach of using log-barriers with the min-max optimization by Moskovitz et al. (2024); Dai et al. (2024) to find the saddle-point of the Lagrangian (MMs). In our experience, learning the Lagrange multipliers in this way is extremely sensitive to the choice of the learning rate. Moreover, to avoid loading an additional reference LLM during fine-tuning, as is done by Moskovitz et al. (2024); Dai et al. (2024), our implementation of MM minimizes the task objective directly (as is done by L3Ms as well).

Fig. 3 shows the achieved task perplexities and helpful-harmless rewards for the different LLMs. At initialization, the helpful-harmless rewards are both low, with a high task perplexity of 1.316 ± 0.4 . This is improved upon by the SFT model, reducing the perplexity to 0.805 ± 0.3 and attaining a high helpfulness reward (due to the task data instilling instruction-following capabilities). Further, MMs sacrifice task performance to over-satisfy the constraints, with mean task perplexities ≥ 0.822 . Conversely, L3Ms satisfy the imposed helpful-harmless reward constraints with consistently lower task perplexities (the mean perplexities are in the range 0.804-0.820). We attribute this to the L3Ms having better Lagrange multipliers by design rather than learning them as in MMs (cf. Section 5.2).




Figure 3: **Helpful and harmless L3Ms.** We report the helpful-harmless rewards and task perplexities achieved by the different LLMs. *Left:* The helpful-harmless rewards attained by the LLM at initialization (at the bottom-left in blue), the SFT model (at the top-left in orange), the MMs (in green), and the L3Ms (in red). We depict the imposed constraints in black, with the dotted gray lines connecting LLMs to their corresponding constraints. Note that constraints are satisfied if the obtained reward point is at the top-right of its corresponding constraint point. *Right:* The mean and standard deviation of the task perplexities for MMs and L3Ms, along with their corresponding constraints; the task perplexity at initialization is $1.316_{\pm 0.4}$ and that of the SFT model is $0.805_{\pm 0.3}$.

7 CONCLUSIONS

In this work, we formulate the SFT and alignment processes for LLMs as one of constrained optimization: we minimize the task objective while simultaneously imposing custom constraints on preferences. This enables customization of the LLM to different preferential properties while maintaining performance on the original task. Consequently, we propose Lagrange Large Language Models (L3Ms) to solve this constrained problem by incorporating the constraints in the objective function using the logarithmic barrier. We include experimental results to illustrate the customization qualities of L3Ms, which can fit to different preferences, providing a personalized user experience.

REFERENCES

- Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. Unified pre-training for program understanding and generation. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (eds.), *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 2655–2668, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.211. URL <https://aclanthology.org/2021.naacl-main.211>. 1
- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. Training a helpful and harmless assistant with reinforcement learning from human feedback, 2022. URL <https://arxiv.org/abs/2204.05862>. 1, 3, 8, 14
- Michiel Bakker, Martin Chadwick, Hannah Sheahan, Michael Tessler, Lucy Campbell-Gillingham, Jan Balaguer, Nat McAleese, Amelia Glaese, John Aslanides, Matt Botvinick, and Christopher Summerfield. Fine-tuning language models to find agreement among humans with diverse prefer-

- 486 ences. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances*
 487 *in Neural Information Processing Systems*, volume 35, pp. 38176–38189. Curran Associates, Inc.,
 488 2022. URL [https://proceedings.neurips.cc/paper_files/paper/2022/](https://proceedings.neurips.cc/paper_files/paper/2022/file/f978c8f3b5f399cae464e85f72e28503-Paper-Conference.pdf)
 489 [file/f978c8f3b5f399cae464e85f72e28503-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/f978c8f3b5f399cae464e85f72e28503-Paper-Conference.pdf). 3
- 490
 491 Ralph Allan Bradley and Milton E. Terry. Rank analysis of incomplete block designs: The method
 492 of paired comparisons. *Biometrika*, 39(3-4):324–345, 12 1952. ISSN 0006-3444. doi: 10.1093/
 493 [biomet/39.3-4.324](https://doi.org/10.1093/biomet/39.3-4.324). URL <https://doi.org/10.1093/biomet/39.3-4.324>. 14
- 494 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhari-
 495 wal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agar-
 496 wal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh,
 497 Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz
 498 Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec
 499 Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In
 500 H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neu-*
 501 *ral Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc.,
 502 2020. URL [https://proceedings.neurips.cc/paper_files/paper/2020/](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6b6fcb4967418bfb8ac142f64a-Paper.pdf)
 503 [file/1457c0d6b6fcb4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6b6fcb4967418bfb8ac142f64a-Paper.pdf). 1
- 504 Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier
 505 Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, Tony Tong Wang,
 506 Samuel Marks, Charbel-Raphael Segerie, Micah Carroll, Andi Peng, Phillip Christoffersen,
 507 Mehul Damani, Stewart Slocum, Usman Anwar, Anand Siththaranjan, Max Nadeau, Eric J
 508 Michaud, Jacob Pfau, Dmitrii Krasheninnikov, Xin Chen, Lauro Langosco, Peter Hase, Erdem
 509 Biyik, Anca Dragan, David Krueger, Dorsa Sadigh, and Dylan Hadfield-Menell. Open problems
 510 and fundamental limitations of reinforcement learning from human feedback. *Transactions on*
 511 *Machine Learning Research*, 2023. ISSN 2835-8856. URL [https://openreview.net/](https://openreview.net/forum?id=bx24KpJ4Eb)
 512 [forum?id=bx24KpJ4Eb](https://openreview.net/forum?id=bx24KpJ4Eb). Survey Certification. 4
- 513 Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario
 514 Amodei. Deep reinforcement learning from human preferences. In I. Guyon, U. Von
 515 Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Ad-*
 516 *vances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.,
 517 2017. URL [https://proceedings.neurips.cc/paper_files/paper/2017/](https://proceedings.neurips.cc/paper_files/paper/2017/file/d5e2c0adad503c91f91df240d0cd4e49-Paper.pdf)
 518 [file/d5e2c0adad503c91f91df240d0cd4e49-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/d5e2c0adad503c91f91df240d0cd4e49-Paper.pdf). 3
- 519 Frank E. Curtis, Vyacheslav Kungurtsev, Daniel P. Robinson, and Qi Wang. A stochastic-gradient-
 520 based interior-point algorithm for solving smooth bound-constrained optimization problems,
 521 2024. URL <https://arxiv.org/abs/2304.14907>. 6
- 522 Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong
 523 Yang. Safe RLHF: Safe reinforcement learning from human feedback. In *The Twelfth Interna-*
 524 *tional Conference on Learning Representations*, 2024. URL [https://openreview.net/](https://openreview.net/forum?id=TyFrPOKYXw)
 525 [forum?id=TyFrPOKYXw](https://openreview.net/forum?id=TyFrPOKYXw). 4, 7, 8
- 526 Ning Ding, Yulin Chen, Bokai Xu, Yujia Qin, Shengding Hu, Zhiyuan Liu, Maosong Sun, and
 527 Bowen Zhou. Enhancing chat language models by scaling high-quality instructional conversa-
 528 tions. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Confer-*
 529 *ence on Empirical Methods in Natural Language Processing*, pp. 3029–3051, Singapore, December
 530 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.emnlp-main.183. URL
 531 <https://aclanthology.org/2023.emnlp-main.183>. 7, 14
- 532
 533 Christian Feller and Christian Ebenbauer. Relaxed logarithmic barrier function based model predic-
 534 tive control of linear systems. *IEEE Transactions on Automatic Control*, 62(3):1223–1238, 2017.
 535 doi: 10.1109/TAC.2016.2582040. 5
- 536 Leo Gao, John Schulman, and Jacob Hilton. Scaling laws for reward model overoptimization. In An-
 537 dreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan
 538 Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume
 539 202 of *Proceedings of Machine Learning Research*, pp. 10835–10866. PMLR, 23–29 Jul 2023.
 URL <https://proceedings.mlr.press/v202/gao23h.html>. 3

- 540 Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. URL
541 <http://www.deeplearningbook.org>. 7
542
- 543 Shane Griffith, Kaushik Subramanian, Jonathan Scholz, Charles L Isbell, and Andrea L
544 Thomaz. Policy shaping: Integrating human feedback with reinforcement learning. In
545 C.J. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger (eds.), *Ad-*
546 *vances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.,
547 2013. URL [https://proceedings.neurips.cc/paper_files/paper/2013/](https://proceedings.neurips.cc/paper_files/paper/2013/file/e034fb6b66aaccl48f445ddfb08da98-Paper.pdf)
548 [file/e034fb6b66aaccl48f445ddfb08da98-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2013/file/e034fb6b66aaccl48f445ddfb08da98-Paper.pdf). 3
- 549 Yiju Guo, Ganqu Cui, Lifan Yuan, Ning Ding, Zexu Sun, Bowen Sun, Huimin Chen, Ruobing Xie,
550 Jie Zhou, Yankai Lin, Zhiyuan Liu, and Maosong Sun. Controllable preference optimization: To-
551 ward controllable multi-objective alignment, 2024. URL [https://arxiv.org/abs/2402.](https://arxiv.org/abs/2402.19085)
552 [19085](https://arxiv.org/abs/2402.19085). 8
- 553 John Hauser and Alessandro Saccon. A barrier function method for the optimization of trajectory
554 functionals with constraints. In *Proceedings of the 45th IEEE Conference on Decision and Con-*
555 *trol*, pp. 864–869, 2006. 5
- 556 Jiaming Ji, Mickel Liu, Josef Dai, Xuehai Pan, Chi Zhang, Ce Bian, Boyuan Chen, Ruiyang
557 Sun, Yizhou Wang, and Yaodong Yang. Beavertails: Towards improved safety align-
558 ment of llm via a human-preference dataset. In A. Oh, T. Naumann, A. Globerson,
559 K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Pro-*
560 *cessing Systems*, volume 36, pp. 24678–24704. Curran Associates, Inc., 2023. URL
561 [https://proceedings.neurips.cc/paper_files/paper/2023/file/](https://proceedings.neurips.cc/paper_files/paper/2023/file/4dbb61cb68671edc4ca3712d70083b9f-Paper-Datasets_and_Benchmarks.pdf)
562 [4dbb61cb68671edc4ca3712d70083b9f-Paper-Datasets_and_Benchmarks.](https://proceedings.neurips.cc/paper_files/paper/2023/file/4dbb61cb68671edc4ca3712d70083b9f-Paper-Datasets_and_Benchmarks.pdf)
563 [pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/4dbb61cb68671edc4ca3712d70083b9f-Paper-Datasets_and_Benchmarks.pdf). 1, 3, 4, 8
- 564 William Karush. Minima of functions of several variables with inequalities as side conditions.
565 Master’s thesis, Department of Mathematics, University of Chicago, Illinois, USA, 1939. 5
566
- 567 Hoel Kervadec, Jose Dolz, Jing Yuan, Christian Desrosiers, Eric Granger, and Ismail Ben Ayed.
568 Constrained deep networks: Lagrangian optimization via log-barrier extensions. In *2022 30th*
569 *European Signal Processing Conference (EUSIPCO)*, pp. 962–966, 2022. doi: 10.23919/
570 [EUSIPCO55093.2022.9909927](https://doi.org/10.23919/EUSIPCO55093.2022.9909927). 5
- 571 W. Bradley Knox and Peter Stone. Tamer: Training an agent manually via evaluative reinforcement.
572 In *2008 7th IEEE International Conference on Development and Learning*, pp. 292–297, 2008.
573 doi: 10.1109/DEVLRN.2008.4640845. 3
574
- 575 Harold W. Kuhn and Albert W. Tucker. Nonlinear programming. In J. Neyman (ed.), *Proceedings*
576 *of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pp. 481–492,
577 Berkeley, California, USA, 1951. University of California Press. 5
- 578 Kaiwen Li, Tao Zhang, and Rui Wang. Deep reinforcement learning for multiobjective optimization.
579 *IEEE Transactions on Cybernetics*, 51(6):3103–3114, 2021. doi: 10.1109/TCYB.2020.2977661.
580 1, 3
- 581 Takuya Makino, Tomoya Iwakura, Hiroya Takamura, and Manabu Okumura. Global optimization
582 under length constraint for neural text summarization. In Anna Korhonen, David Traum, and
583 Lluís Màrquez (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computa-*
584 *tional Linguistics*, pp. 1039–1048, Florence, Italy, July 2019. Association for Computational Lin-
585 guistics. doi: 10.18653/v1/P19-1099. URL <https://aclanthology.org/P19-1099>.
586 7
- 587 Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih. Monte carlo gradient esti-
588 mation in machine learning. *Journal of Machine Learning Research*, 21(132):1–62, 2020. URL
589 <http://jmlr.org/papers/v21/19-346.html>. 7
590
- 591 Michael Moor, Oishi Banerjee, Zahra Shakeri Hossein Abad, Harlan M. Krumholz, Jure Leskovec,
592 Eric J. Topol, and Pranav Rajpurkar. Foundation models for generalist medical artificial
593 intelligence. *Nature*, 616(7956):259–265, Apr 2023. ISSN 1476-4687. doi: 10.1038/
[s41586-023-05881-4](https://doi.org/10.1038/s41586-023-05881-4). URL <https://doi.org/10.1038/s41586-023-05881-4>. 1

- 594 Ted Moskovitz, Aaditya K Singh, DJ Strouse, Tuomas Sandholm, Ruslan Salakhutdinov, Anca
595 Dragan, and Stephen Marcus McAleer. Confronting reward model overoptimization with con-
596 strained RLHF. In *The Twelfth International Conference on Learning Representations*, 2024.
597 URL <https://openreview.net/forum?id=gkfvvn0fLU>. 4, 7, 8
- 598
599 Stephen G. Nash, R. Polyak, and Ariela Sofer. *A Numerical Comparison of Barrier and Modi-*
600 *fied Barrier Methods For Large-Scale Bound-Constrained Optimization*, pp. 319–338. Springer
601 US, Boston, MA, 1994. ISBN 978-1-4613-3632-7. doi: 10.1007/978-1-4613-3632-7_16. URL
602 https://doi.org/10.1007/978-1-4613-3632-7_16. 5
- 603 OpenAI. Gpt-4 technical report, 2024. URL <https://arxiv.org/abs/2303.08774>. 1
- 604
605 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
606 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kel-
607 ton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike,
608 and Ryan Lowe. Training language models to follow instructions with human feedback. In
609 S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in*
610 *Neural Information Processing Systems*, volume 35, pp. 27730–27744. Curran Associates, Inc.,
611 2022. URL [https://proceedings.neurips.cc/paper_files/paper/2022/](https://proceedings.neurips.cc/paper_files/paper/2022/file/blfede53be364a73914f58805a001731-Paper-Conference.pdf)
612 [file/blfede53be364a73914f58805a001731-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/blfede53be364a73914f58805a001731-Paper-Conference.pdf). 2
- 613
614 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
615 Finn. Direct preference optimization: Your language model is secretly a reward model. In
616 A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in*
617 *Neural Information Processing Systems*, volume 36, pp. 53728–53741. Curran Associates, Inc.,
618 2023. URL [https://proceedings.neurips.cc/paper_files/paper/2023/](https://proceedings.neurips.cc/paper_files/paper/2023/file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf)
619 [file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf). 14
- 620
621 Alexandre Rame, Guillaume Couairon, Corentin Dancette, Jean-Baptiste Gaya, Mustafa
622 Shukor, Laure Soulier, and Matthieu Cord. Rewarded soups: towards pareto-optimal
623 alignment by interpolating weights fine-tuned on diverse rewards. In A. Oh, T. Nau-
624 mann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural*
625 *Information Processing Systems*, volume 36, pp. 71095–71134. Curran Associates, Inc.,
626 2023. URL [https://proceedings.neurips.cc/paper_files/paper/2023/](https://proceedings.neurips.cc/paper_files/paper/2023/file/e12a3b98b67e8395f639fde4c2b03168-Paper-Conference.pdf)
627 [file/e12a3b98b67e8395f639fde4c2b03168-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/e12a3b98b67e8395f639fde4c2b03168-Paper-Conference.pdf). 1, 3, 4
- 628
629 Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi
630 Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Ev-
631 timov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafori, Wenhan Xiong,
632 Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier,
633 Thomas Scialom, and Gabriel Synnaeve. Code llama: Open foundation models for code, 2024.
634 URL <https://arxiv.org/abs/2308.12950>. 1
- 635
636 John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-
637 dimensional continuous control using generalized advantage estimation. In *Proceedings of the*
638 *International Conference on Learning Representations (ICLR)*, 2016. 6, 15
- 639
640 A Ben Tal, M Tsibulevskii, and I Yusefovich. Modified barrier methods for constrained and minimax
641 problems. *Technical report, Optimization Laboratory, Israel Institute of Technology*, 1992. 5
- 642
643 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
644 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Ar-
645 mand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation
646 language models, 2023. URL <https://arxiv.org/abs/2302.13971>. 7
- 647
648 Yue Wang, Weishi Wang, Shafiq Joty, and Steven C.H. Hoi. CodeT5: Identifier-aware unified
649 pre-trained encoder-decoder models for code understanding and generation. In Marie-Francine
650 Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021*
651 *Conference on Empirical Methods in Natural Language Processing*, pp. 8696–8708, Online
652 and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguis-
653 tics. doi: 10.18653/v1/2021.emnlp-main.685. URL [https://aclanthology.org/2021.](https://aclanthology.org/2021.emnlp-main.685)
654 [emnlp-main.685](https://aclanthology.org/2021.emnlp-main.685). 1

- 648 Zihao Wang, Chirag Nagpal, Jonathan Berant, Jacob Eisenstein, Alexander Nicholas D’Amour,
649 Sanmi Koyejo, and Victor Veitch. Transforming and combining rewards for aligning large lan-
650 guage models. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria
651 Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International
652 Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*,
653 pp. 51161–51176. PMLR, 21–27 Jul 2024. URL [https://proceedings.mlr.press/
654 v235/wang24ay.html](https://proceedings.mlr.press/v235/wang24ay.html). 8
- 655 Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi,
656 Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick
657 von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger,
658 Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-art natural
659 language processing. In Qun Liu and David Schlangen (eds.), *Proceedings of the 2020 Confer-
660 ence on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–
661 45, Online, October 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.
662 emnlp-demos.6. URL <https://aclanthology.org/2020.emnlp-demos.6.7.14>
- 663 Zeqiu Wu, Yushi Hu, Weijia Shi, Nouha Dziri, Alane Suhr, Prithviraj Ammanabrolu,
664 Noah A Smith, Mari Ostendorf, and Hannaneh Hajishirzi. Fine-grained human feed-
665 back gives better rewards for language model training. In A. Oh, T. Naumann,
666 A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural In-
667 formation Processing Systems*, volume 36, pp. 59008–59033. Curran Associates, Inc.,
668 2023. URL [https://proceedings.neurips.cc/paper_files/paper/2023/
669 file/b8c90b65739ae8417e61eadb521f63d5-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/b8c90b65739ae8417e61eadb521f63d5-Paper-Conference.pdf). 1, 3, 4
- 670 Xi Yang, Aokun Chen, Nima PourNejatian, Hoo Chang Shin, Kaleb E. Smith, Christopher Parisien,
671 Colin Compas, Cheryl Martin, Anthony B. Costa, Mona G. Flores, Ying Zhang, Tanja Magoc,
672 Christopher A. Harle, Gloria Lipori, Duane A. Mitchell, William R. Hogan, Elizabeth A.
673 Shenkman, Jiang Bian, and Yonghui Wu. A large language model for electronic health records.
674 *npj Digital Medicine*, 5(1):194, Dec 2022. ISSN 2398-6352. doi: 10.1038/s41746-022-00742-2.
675 URL <https://doi.org/10.1038/s41746-022-00742-2>. 1
- 676
677 Jingzhao Zhang, Tianxing He, Suvrit Sra, and Ali Jadbabaie. Why gradient clipping accelerates
678 training: A theoretical justification for adaptivity. In *International Conference on Learning Rep-
679 resentations*, 2020. URL <https://openreview.net/forum?id=BJgnXpVYwS>. 7
- 680 Zhanhui Zhou, Jie Liu, Jing Shao, Xiangyu Yue, Chao Yang, Wanli Ouyang, and Yu Qiao. Be-
681 yond one-preference-fits-all alignment: Multi-objective direct preference optimization. In Lun-
682 Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computa-
683 tional Linguistics ACL 2024*, pp. 10586–10613, Bangkok, Thailand and virtual meeting, August
684 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.630. URL
685 <https://aclanthology.org/2024.findings-acl.630.4.8>
- 686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

702 A EXPERIMENTAL SETUP

703
704
705 In addition to Section 6.1, here we provide further details of the experimental setup.

706
707 **Task data** We use UltraChat (Ding et al., 2023), a large-scale dataset of instructional conversa-
708 tions, as our task data to induce instruction-following capabilities. Since each sample contains a
709 sequence of multi-turn question-answer pairs, we randomly sample one of the answers as the re-
710 sponse and treat the preceding dialogue as the prompt. We then filter such (prompt, response) pairs
711 to a maximum token length of 512. Consequently, we obtain 340k training samples, 1.7k validation
712 samples, and 1.7 test samples, split randomly as the dataset does not contain train-val-test splits.

713
714 **Hyper-parameters** We fine-tune LLMs for 1 epoch on the task data, with a mini-batch size of
715 64. We use Adam with a learning rate of 10^{-6} , and a cosine learning rate scheduler (with 5% of the
716 epoch used for warmup). We set weight decay to 0.1 and the gradient clipping maximum norm to 1.
717 We utilize 16-bit (mixed) precision training and gradient checkpointing. We exponentially decay the
718 log-barrier parameter μ during fine-tuning from 1 to 10^{-6} , and use a smoothing factor of 0.1 for the
719 exponential moving average. Lastly, we use top- p sampling (p set to 0.9) for response generation.
720 Apart from this, we use the default hyper-parameters in the Transformers library (Wolf et al., 2020).

721 A.1 LEARNING PREFERENCE REWARD MODELS

722
723 While some preference reward functions are engineered or rule-based, others are learned. Such
724 preferences can often be difficult to quantify. Alternatively, it is easier to compare responses with
725 respect to the preference, e.g., ranking them from most to least helpful. Consequently, the data for
726 learning preference reward models consist of tuples of the form (x, y_+, y_-) , where the prompt x is
727 accompanied by two responses y_+ and y_- , with the former response being preferred over the latter.

728 The preference reward model is denoted by $r_\phi(\cdot)$ (parameterized by ϕ). Assuming the Bradley-Terry
729 model (Bradley & Terry, 1952), the model’s predicted probability for preferring y_+ over y_- is:

$$730 p_{r_\phi}(y_+ \succ y_- | x) = \sigma(r_\phi(y_+ | x) - r_\phi(y_- | x)),$$

731 with the standard logistic function $\sigma(\cdot)$. Then, the model minimizes the negative log-likelihood:

$$732 \min_{\phi} \mathbb{E}_{(x, y_+, y_-) \sim t(\cdot)} [-\log p_{r_\phi}(y_+ \succ y_- | x)].$$

733
734 Taking inspiration from Rafailov et al. (2023), we initialize the preference reward model $r_\phi(\cdot)$ as
735 a pre-trained LLM, and set the reward to be its length-normalized log-likelihood. In this way, we
736 utilize the pre-trained model fully, not just its backbone. As the preference reward model is fine-
737 tuned, its log-likelihoods are updated to differentiate the preferred responses from the rejected ones.

738
739 **Helpful and harmless data** We use the Helpful and Harmless (Bai et al., 2022) preference data
740 to learn two reward models respectively. We obtain 161k training samples and 9k test samples after
741 filtering the (prompt, response) pairs to a maximum token length of 2024; 3/4-th are for helpfulness
742 while the remaining 1/4-th are for harmlessness. We further set 5% of the training data for validation.

743
744 **Hyper-parameters** We initialize all preference reward functions with LLaMA-7B. We fine-tune
745 for 2 epochs with a mini-batch size of 64. We use Adam with a learning rate of 10^{-6} , and a cosine
746 learning rate scheduler (with 10% of the epoch used for warmup). We set weight decay to 0.1 and
747 the gradient clipping maximum norm to 1. We utilize 16-bit (mixed) precision training and gradient
748 checkpointing. Apart from this, we use the default hyper-parameters in the Transformers library
749 (Wolf et al., 2020). We validate after every 10% of the epoch and save the best model.

756 B POLICY GRADIENT

757
758 We are interested in the policy gradient $\partial_\theta C_i(\theta)$. Note that when taking derivatives with respect to
759 the parameters of a distribution in an expectation, we can use the log-derivative trick:
760

$$761 \quad \partial_\theta \mathbb{E}_{x \sim p_\theta(\cdot)} [f(x)] = \int dx f(x) \partial_\theta p_\theta(x) = \int dx f(x) \frac{p_\theta(x)}{p_\theta(x)} \partial_\theta p_\theta(x)$$

$$762 \quad = \int dx f(x) p_\theta(x) \partial_\theta \log p_\theta(x) = \mathbb{E}_{x \sim p_\theta(\cdot)} [f(x) \partial_\theta \log p_\theta(x)].$$

763
764
765 Applying this to the policy gradient $\partial_\theta C_i(\theta)$ yields,
766

$$767 \quad \partial_\theta C_i(\theta) = \partial_\theta \mathbb{E}_{\substack{(x,\cdot) \sim p(\cdot) \\ y \sim \pi_\theta(\cdot|x)}} [c_i(y|x)] = \mathbb{E}_{\substack{(x,\cdot) \sim p(\cdot) \\ y \sim \pi_\theta(\cdot|x)}} [c_i(y|x) \partial_\theta \log \pi_\theta(y|x)], \quad (12)$$

768
769 where $c_i(y|x) = b_i - r_i(y|x)$. This is the simplest form of the policy gradient and can be estimated as
770 Monte Carlo averages. We refer readers to Schulman et al. (2016) for a review of other estimates.
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809