

Towards Enhanced RAC Accessibility: Leveraging Datasets and LLMs

Bejarano Sepulveda Edison,

Universidad de Barcelona
ejbejaranos@gmail.com

Potes Hector Nicolai,

nicolai.potes00@gmail.com

Pineda Montoya Santiago,

Universidad Nacional de Colombia
sapinedamo@unal.edu.co

Rodriguez Felipe Ivan,

Fund. Univ. Los Libertadores
ifrodriguez@libertadores.edu.co

Orduy Jaime Enrique,

Fund. Univ. Los Libertadores
jaime.orduy@libertadores.edu.co

Rosales Cabezas Alec,

Fund. Univ. Los Libertadores
amrosalesc@libertadores.edu.co

Traslaviña Navarrete Danny

Fund. Univ. Los Libertadores
dstraslavinan@libertadores.edu.co

Madrid Farfan Sergio,

Fund. Univ. Los Libertadores
snmadridf@libertadores.edu.co

Abstract

This paper explores the potential of large language models (LLMs) to make the Aeronautical Regulations of Colombia (RAC) more accessible. Given the complexity and extensive technicality of the RAC, this study introduces a novel approach to simplifying these regulations for broader understanding. By developing the first-ever RAC database, which contains 24,478 expertly labeled question-and-answer pairs, and fine-tuning LLMs specifically for RAC applications, the paper outlines the methodology for dataset assembly, expert-led annotation, and model training. Utilizing the Gemma1.1 2b model along with advanced techniques like Unsloth for efficient VRAM usage and flash attention mechanisms, the research aims to expedite training processes. This initiative establishes a foundation to enhance the comprehensibility and accessibility of RAC, potentially benefiting novices and reducing dependence on expert consultations for navigating the aviation industry's regulatory landscape.

You can visit the [dataset](#) and the [model](#) here.

1 Introduction

The Colombian aviation industry operates under the Aeronautical Regulations of Colombia (RAC) ([Aerocivil Colombia, 2023](#)), a comprehensive legal framework that comprises approximately 50 detailed regulations and manuals. It is worth noting that the RAC is currently undergoing harmonization with the Latin American Aeronautical Regulations (LAR). The technical complexity and voluminous nature of these documents pose significant challenges to accessibility. However, the advent of LLMs promises to revolutionize this scenario by simplifying complex texts, making regulatory information more understandable and accessible

to a broader audience. As highlighted by Yang et al. ([Yang et al., 2023](#)), LLMs such as ChatGPT have demonstrated significant potential in various practical applications, suggesting their utility in interpreting and simplifying legal and regulatory texts. By translating legal terminology into plain language, LLMs play a crucial role in demystifying aviation regulations, thereby improving understanding and compliance within the industry.

2 Objectives

The overarching goal is to utilize LLMs to make the Aeronautical Regulations of Colombia (RAC) more accessible and understandable for both aviation professionals and the general public. The approach is multifaceted: First, we aim to develop a comprehensive dataset from the RAC's initial five documents, laying the groundwork for LLM training. Second, this project seeks collaboration with industry and academic experts to annotate and refine this dataset, ensuring its relevance and accuracy. Third, we will train LLMs with the curated dataset, and lastly, evaluate their performance in simplifying the RAC's content based on feedback from aeronautical experts, thereby enhancing regulatory compliance and understanding.

3 Background

Innovative tech solutions are crucial for navigating the complexities of the Aeronautical Regulations of Colombia (RAC). Starting with decision trees to enhance access to RAC, the aviation sector's push for safety and efficiency demands more sophisticated approaches, given its high reliability standards ([Rodríguez-Palomino, 2021](#); [Abubakar et al., 2022](#)). Early AI uses in aviation, such

as expert systems for decision support, highlight AI’s impact on safety and operational efficiency (EAS, 2020). The emergence of LLMs, including Seq2Seq and Transformer architectures, marks a significant advancement, offering detailed conversational AI support (MDP, 2021; CMS, 2021; IEE, 2021). A UAEAC (Unidad Administrativa Especial de Aeronáutica Civil) project demonstrated AI’s potential for real-time RAC inquiries, showcasing AI’s role in enhancing regulatory consultations (Rodriguez, 2021). This move towards LLMs for regulatory compliance and decision support illustrates a shift to intelligent solutions, with LLMs providing the flexibility and depth to address the regulatory domain’s complexities effectively.

4 Methodology

This research systematically applies LLMs to enhance the accessibility of the Aeronautical Regulations of Colombia (RAC). Our methodology includes dataset generation, expert-driven labeling, and iterative LLM fine-tuning, ensuring a comprehensive, data-driven approach.

4.1 Dataset Generation from RAC Documents

The dataset was crafted from the RAC using an automated process, as depicted in Figure 1. This process began with converting PDFs to text, followed by processing through a GPT API or a similar model. The system handled two pages at a time and iteratively compiled the extracted data. The resulting dataset, comprising questions, answers, and relevant RAC references, was thus assembled for further analysis.

4.2 Labeling Process for the RAC Dataset

The RAC Dataset underwent refinement using the Argilla framework within the Hugging Face environment, as illustrated in Figure 2. This tool facilitated structured annotation tasks, leveraging the expertise of aeronautical engineering specialists from Fundación Universitaria Los Libertadores. They assessed each sample for quality, retaining those ranked above 3 and removing lower-ranked ones, thus ensuring the dataset’s integrity.

The concluding phase of the process involved removing the discarded samples. Consequently, a high-quality dataset was consolidated and made available through Hugging Face.

4.3 Fine tuning models

In this study, we fine-tuned the pre-trained GEMMA model for a specific NLP task by employing a Parameter-Efficient Fine-Tuning (PEFT) strategy integrated with Low-Rank Adaptation (LoRA). Initially, the somosnlp/ColombiaRACFullyCurated dataset was tokenized and divided into training and testing subsets.

Version	Data Quality	Model Type
V8	Clean	gemma-1.1-2b-it
V7	Clean	gemma-1.1-2b-it
V6	50% Clean	gemma-2b-it
V5	50% Clean	gemma-2b-it
V4	50% Clean	gemma-2b-it
V3	Raw	gemma-2b-it
V2	Raw	gemma-2b-it
V1	Raw	gemma-2b-it

Table 1: Model Versions, Quality, and Types

The model was then configured to efficiently adapt to the task by selectively modifying a minimal subset of its parameters, utilizing PEFT and LoRA techniques. Key hyperparameters for the training included a learning rate of 5e-5 and a batch size of 3, optimized using the AdamW optimizer with a weight decay of 0.001.

Version	Techniques	GPU Type
V8	flashattn+perft+qlora	RTX 3090
V7	unsloth+peft+qlora	L4-24GB
V6	unsloth+peft+qlora	L4-24GB
V5	unsloth+peft+qlora	A100-40GB
V4	peft+qlora	A100-40GB
V3	peft+qlora	A100-40GB
V2	peft+qlora	A100-40GB
V1	peft+qlora	A100-40GB

Table 2: Summary of Techniques and GPU Type

The training, managed by the SFTTrainer, emphasized gradient accumulation and learning rate scheduling to enhance model performance. After training, the fine-tuned model was combined with LoRA weights for deployment. This process demonstrates an efficient approach to customizing large-scale language models for specific tasks, leveraging advanced techniques to balance computational efficiency and model effectiveness. PyTorch and Hugging Face transformers were utilized for implementation due to their robustness and support for complex NLP tasks.

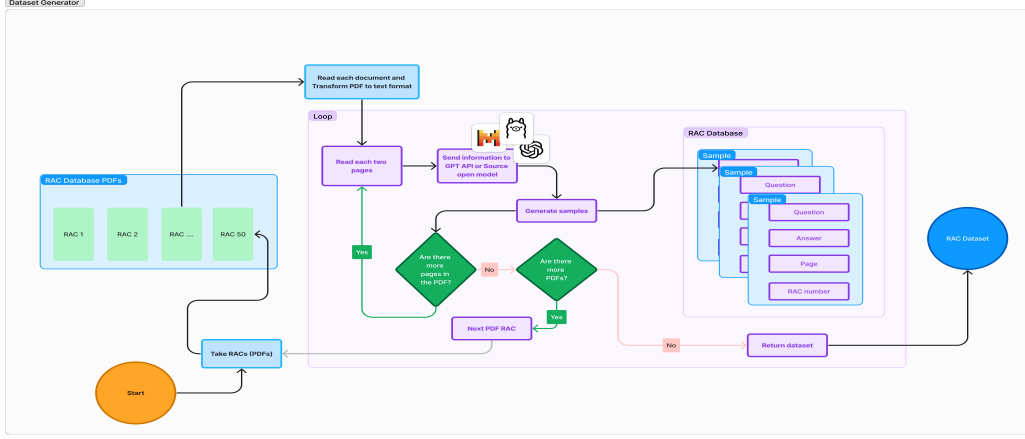


Figure 1: Process for data extraction from RAC PDFs using GPT API.

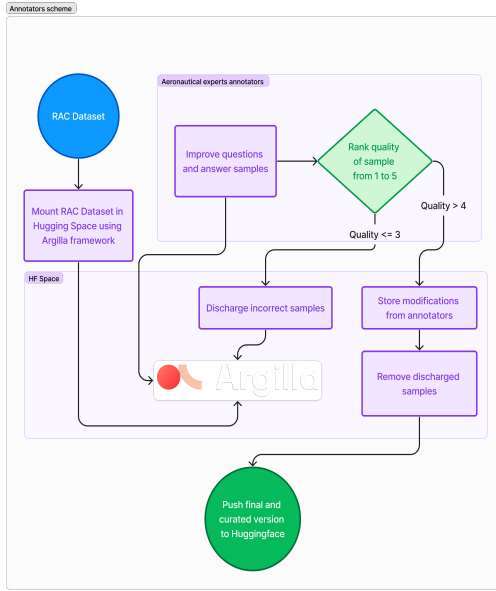


Figure 2: Flow diagram for system annotator.

	Time(s)	Loss	FLOPs
V8	12,607	0.194	3.94E+14
V7	33,262	0.194	2.51E+15
V6	1,977	0.243	5.01E+15
V5	1,779	0.076	5.18E+15
V4	1,833	0.092	5.06E+15
V3	3,987	0.071	4.99E+15
V2	4,239	1.241	4.95E+15
V1	50,973	0.6	2.29E+15

Table 3: Summary of Results: Runtime, Loss, FLOPs

areas needing improvement, while high ratings in RACs 1 and 5 suggest strengths. These results confirm the model’s potential for accuracy and generalization, though RAC 3 requires adjustments.

RAC	Mean	Median	Number of Tests
1	6.12	7	100
2	5.8	6.5	20
3	3.464	1	56
5	6.34	7.5	100
General	5.6	7	276

Table 4: Expert Evaluations by RAC: Mean, Median, and Test Count

5 Results

5.1 Quantitative results

The quantitative results (see Table 3) highlight the effectiveness of the efforts. Earlier versions (V5 to V3) had low loss but lacked answer quality, while subsequent iterations improved response accuracy, training efficiency, and environmental sustainability. V8 was the most optimized model, enhancing RAC interpretability with advanced LLMs, lower training loss, and reduced environmental impact.

5.2 Qualitative results

Table 4 shows the model’s strong performance with average scores of 7 from 276 tests. However, RAC 3’s low scores (mean 3.464, median 1) indicate

6 Conclusions

The overarching goal was to enhance accessibility to the RAC through the utilization of LLMs. This objective was pursued through a multi-faceted approach: first, by developing a comprehensive dataset from the initial RAC documents; second, by collaborating with industry and academic experts to refine the dataset; third, by training LLMs with this curated dataset; and finally, by evaluating their performance. This initiative aims to improve regulatory compliance and understanding among aviation professionals and the general public.

References

2020. [Easa ai roadmap v1.0](#).
2021. [Advances in conversational agents](#). *MDPI Electronics*, 13(8):1409.
2021. [Latest trends in conversational agents](#). In *CMS Conferences*.
2021. [Safari: A new age of conversational agents](#). *IEEE Explore*.
- Mahmoud Abubakar, Odunlami EriOluwa, Mangai Teyei, and Fadi Al-Turjman. 2022. Ai application in the aviation sector. In *2022 International Conference on Artificial Intelligence of Things and Crowdsensing (AIoTCs)*, pages 52–55. IEEE.
- Aerocivil Colombia. 2023. Reglamentos aeronáuticos de colombia. <https://www.aerocivil.gov.co/autoridad-de-la-aviacion-civil/reglamentacion/rac>. Aeronautical Regulations of Colombia.
- Paulino Rodriguez. 2021. [Artificial intelligence for the administration of the aeronautical regulations of colombia \(rac\)](#). *UNAD Journal*.
- Paulino Rodríguez-Palomino. 2021. [Inteligencia artificial para la administración de los reglamentos aeronáuticos de colombia \(rac\)](#). *Revista Estrategia Organizacional*, 10(1). Recibido: 8 de agosto de 2020; Evaluado: 7 de septiembre de 2020; Aprobado: 14 de noviembre de 2020.
- Jingfeng Yang, Hongye Jin, Ruixiang Tang, Xiaotian Han, Qizhang Feng, Haoming Jiang, Bing Yin, and Xia Hu. 2023. [Harnessing the power of llms in practice: A survey on chatgpt and beyond](#).