Ryan Wallace

Boston, MA | ryan@ryancswallace.dev | 617-852-9239
github.com/ryancswallace | **ryancswallace.dev** | linkedin.com/in/ryancswallace

October 2, 2026

Hiring Team
Paperless Parts

Re: Staff Machine Learning Engineer

Dear Hiring Team,

I am a Lead Data Scientist at the Federal Reserve Bank of Boston. I am applying for the Staff Machine Learning Engineer role at Paperless Parts because I am drawn to applying ML to the drawings and documents behind real-world manufacturing, and your emphasis on rigorous evaluation, practical model choices, and ownership from prototype through production closely matches how I approach building AI systems. I combine applied modeling with the software engineering needed to deploy and maintain reliable ML systems.

I built regularly used classifiers for debit-card transaction descriptions that saved researchers hundreds of hours of manual labeling. I trained random forests on text embeddings and LightGBM using human annotations and LLM-generated labels produced through constrained decoding. I also built and served a retrieval-augmented generation application with vLLM so researchers could analyze internal reports using information available as of specified dates.

I bring data and compute infrastructure experience as well. I architected an AWS lakehouse supporting more than 200 researchers and over 10 PB of data, and helped implement a migration that reduced aggregate data platform costs by 75%. My Spark and Apache Iceberg streaming work cut update latency from roughly a day to seconds and reduced pipeline failures by 80% through improved schema handling. I also designed an NVIDIA H100 GPU cluster and used Hugging Face Transformers for inference and LoRA fine-tuning, evaluating model outputs with human review and LLM-as-a-judge.

Evaluation and failure handling are central to my work. In my open-source Jobman Diagnose project, I combined deterministic diagnostics with optional, schema-validated LLM suggestions, preserving local findings when providers fail or suggestions fail validation. I evaluated it on 72 cases across 216 model executions. For cases requiring a concrete cause, at least 95% of executions produced diagnoses that passed checks for failure classification, required facts, and causal relationships. No unsupported claims were observed in that evaluation. I also built benchmatrix for performance regression testing using repeated measurements and paired statistical comparisons.

In my current role, I mentor two colleagues, run design reviews, and direct maintenance and incident response. I would bring this combination of modeling, systems engineering, and technical leadership to Paperless Parts's work on turning complex manufacturing data into reliable intelligence through repeatable training pipelines, efficient inference services, and benchmarks that make quality and cost trade-offs explicit. Thank you for considering my application; I would welcome a conversation about your team's needs.

Sincerely,
Ryan Wallace