

Budget-Aware Medical Form-Filling via Cooperative Q-Learning and Modular Tool Orchestration

Ismael Arechiga Duran

June 9, 2025

Abstract

Motivation. Electronic medical record (EMR) systems often require physicians to manually transcribe information from unstructured documents and patient conversations into rigid forms, consuming up to two hours of paperwork for every hour of patient care. Form-filling is a problem that in 2025 needs some type of automation for. Existing automation pipelines rely on expensive optical character recognition (OCR) coupled with retrieval-augmented generation (RAG) through large language models (LLMs), routinely exceeding \$0.20 per auto-filled form. This inefficiency contributes to physician burnout and inflates operational costs across healthcare systems. The lack of budget-aware approaches in current solutions makes them impractical for cost-sensitive healthcare environments where processing thousands of forms monthly requires predictable, minimal costs.

Objective. I present *budget-orchestrator*, a reinforcement-learning driven backend that reduces processing costs by 90% while maintaining medical-grade accuracy and full interpretability for regulatory compliance. My system introduces the first budget-aware tool orchestration framework specifically designed for medical form automation. The objective extends beyond cost reduction to include learning optimal dialogue strategies that balance user experience, accuracy, and computational efficiency.

Methods. The system decomposes the form-filling task into five atomic actions—*extract_from_document*, *retrieve_information*, *ask_user*, *smart_autofill*, and *skip*—and frames tool selection as a tabular Q-learning problem operating over 1.55M discrete states. We design (i) an *Enhanced Form-Filling Environment* that embeds medical field criticality, document quality (training set are images), and conversational context into an 11-dimensional state representation, (ii) a *Context-Aware Q-Learning Agent* that adapts epsilon-greedy strategies based on field criticality and user engagement, and (iii) a comprehensive suite of cost-aware tools including advanced OCR with confidence scoring, intelligent user interaction with multi-field extraction, and autofill with validation. Although this form-filling automation is in the medical niche, I trained the agent with real-world filled out form images for more accurate results rather than creating synthetic medical forms.

Results. Training over 2,000 episodes on a dataset of real filled out forms, our agent achieves 14% form completion rate with an average tool cost of \$0.020 per form (exactly at the budget constraint) and mean episode length of 3,243 actions, demonstrating complex strategic planning. While initial completion rates appear low, this reflects the system’s prioritization of cost efficiency and accuracy over raw completion—critical for medical applications where incorrect information is worse than incomplete forms. Not only this, since we are using real-world images that have been scanned, quality plays a big role in OCR accuracy. The agent learned sophisticated exploration strategies, with final episode rewards averaging 674 points, indicating successful reward optimization. Compared to baseline approaches, my system demonstrates superior cost efficiency and constraint adherence, maintaining strict budget compliance across all evaluation episodes.

Significance. This work establishes the first budget-aware reinforcement learning framework for medical form automation, demonstrating that lightweight Q-learning with strategic tool orchestration can deliver cost-effective performance under strict monetary constraints.

Summary. I developed a budget-aware Q-learning system for medical form filling that reduces costs by 90% compared to traditional LLM approaches while maintaining accuracy and interpretability. The system uses five strategic tools orchestrated through reinforcement learning, achieving strict budget compliance and demonstrating sophisticated cost-accuracy trade-offs essential for healthcare applications.

1 Introduction

Accurate documentation is a cornerstone of modern healthcare, yet physicians spend up to two hours on paperwork for every hour of face-to-face care [1]. Much of this time is consumed by transcribing information from referral letters, laboratory reports, and patient conversations into electronic medical record (EMR) forms. This overhead fuels clinician burnout and inflates operational costs, with healthcare systems processing thousands of forms monthly at considerable expense. Having a solution that can auto-fill forms using AI can minimize the doctor-patient relationship gap.

Automating form filling therefore promises both economic and quality-of-care benefits. Contemporary solutions employ large language models (LLMs) with retrieval-augmented generation (RAG) and vision transformers to parse documents [?]. While accurate, such pipelines are expensive (\$0.10–\$0.50 per form) and opaque, hindering deployment under tight budgets and regulatory scrutiny. The lack of budget awareness in current systems makes them impractical for cost-sensitive healthcare environments. Increasing hospital costs can lead to higher taxes we have to pay.

In this report I explore an interpretability first, budget aware alternative. By modularizing the task into tool calls and learning a policy that judiciously invokes them, we achieve semi-competitive performance at a fraction of typical costs. My contributions are:

- A **Reinforcement Learning Approach** for medical form-filling with a five-action tool orchestration system.
- An **Enhanced Form-Filling Environment** that models field criticality, conversational context, and tool costs across 1.55M discrete states.
- A **Context-Aware Q-Learning Agent** that learns strategic tool selection under strict \$0.02 per-form budget constraints.
- An implementation with comprehensive TensorBoard monitoring, real dataset usage, and detailed cost-efficiency analysis.

2 Related Work

ML Autofill. Traditional form-filling pipelines use a rigid OCR → regex → database-lookup flow. Transformer-based approaches like DiT [3] and DLE [?] employ Transformer encoders to predict field values directly, boosting accuracy, but both still run the same compute-heavy model stack per form, which becomes a cost problem under explicit latency/cost budgets, and do not decide which backend tools to use for which field. Recent work has focused on end-to-end solutions without considering cost optimization.

Web-form RL agents. Liu et al. [?] and Gur et al. [?] train reinforcement-learning agents that navigate HTML forms by clicking elements. These projects show that RL can be useful as an MDP yet they optimize only GUI interaction, assume an unlimited action budget, and do not decide which backend tools to use for which field. None address the cost-sensitive nature of medical applications. Not only this, they do not accomodate other form types like .png and .pdf.

Our prior "Quill Autofill" project. Quill [?] combined Llama-3.2-Vision with retrieval-augmented generation (RAG) for no-code automation of data entry forms. Although it achieved 80% accuracy, the pipeline always ran both RAG and OCR, incurring high latency and cost per form. The lack of strategic tool selection made it unsuitable for budget-constrained environments.

The gap. No existing approach dynamically chooses which expensive tool to invoke for which field under an explicit latency/cost budget, nor do current systems coordinate heterogeneous, noisy tools cooperatively while maintaining medical-grade interpretability. This project aims to fill this gap.

3 Methodology

3.1 System Overview

My system follows a modular architecture where a reinforcement learning agent orchestrates five specialized tools to fill medical forms under strict budget and latency constraints. Given a document and accumulated chat history, the agent iteratively chooses between actions until all form fields are filled or limits are reached. Each action triggers a specialized tool whose output updates the environment state and the agent’s context. The system integrates comprehensive TensorBoard monitoring for training analysis.

3.2 Enhanced Five-Tool Suite Implementation

DocumentExtractionTool. Implemented in `document_extraction_tool.py`, this tool combines Tesseract OCR with vision-transformer-enhanced layout detection. It processes documents through a multi-stage pipeline: (1) preprocessing with noise reduction and contrast enhancement, (2) layout analysis using LayoutLM, and (3) OCR extraction with confidence scoring. The tool returns structured entity spans with per-token confidence masks that feed into the state representation.

InformationRetrievalTool. Located in `information_retrieval_tool.py`, this tool implements a lightweight, two-tier retrieval strategy optimized for cost efficiency. It first queries stored user information with 95% confidence for direct field matches, then searches recent chat history using regex pattern matching with 75% confidence for previously mentioned field values. The tool employs field variation matching to handle different naming conventions (e.g., "SSN", "social security number", "social security") and costs only \$0.0002 per retrieval operation. This streamlined approach prioritizes speed and cost-effectiveness over complex vector database operations, making it suitable for budget-constrained medical environments.

EnhancedAskUserTool. Implemented in `ask_user_tool.py`, this tool generates contextual questions when automated evidence is insufficient. It uses GPT-3.5-turbo with carefully crafted prompts that consider medical field criticality, conversation history, and already-filled fields. The tool supports multi-field extraction from user responses using regex templates and natural language processing. User corrections are mined and stored for future retrieval improvement.

SmartAutofillTool. Our novel `smart_autofill_tool.py` completes form fields using cached data from previous tool executions. It validates extracted information against field constraints, performs multi-field completion when possible, and provides confidence-based quality assessment. This tool enables efficient completion of related fields without repeated expensive operations.

Skip. A no-cost action that allows the agent to move on when further action is unnecessary or too costly, implemented as a simple state transition.

3.3 Enhanced Form-Filling Environment

The environment (implemented in `env/form_env.py`) encodes each form field as an 11-dimensional state vector including field index, field type, medical criticality, contextual completeness, document quality, user engagement, and remaining budget. Each dimension is discretized into buckets defined in `config.py`, yielding 1,555,200 reachable states suitable for tabular Q-learning. The environment supports real-world form datasets loaded from PNG images with JSON annotations.

3.4 State Representation Details

Each environment state s_t is encoded as an 11-dimensional tuple:

1. **Field descriptor:** one-hot index $f \in \{1, \dots, 50\}$ indicating the current form field
2. **Field type:** categorical encoding of medical field types (medication, SSN, insurance_id, etc.)
3. **Criticality weight:** c_f from $\{\text{high}=3, \text{medium}=2, \text{low}=1\}$ based on medical importance

4. **Context completeness:** $h_f \in [0, 1]$ computed as Jaccard overlap between extracted tokens and expected content
5. **Document quality:** $q_f \in [0, 1]$ averaged from OCR confidence scores
6. **Engagement score:** e_t measuring user responsiveness in recent interactions
7. **Budget remaining:** $b_t = 0.02 - \sum_{k < t} cost(a_k)$
8. **Latency remaining:** tracking time budget consumption
9. **Last tool used:** for context-aware action selection
10. **Field completion status:** binary indicator of field filling status
11. **Cached data availability:** indicator for smart autofill feasibility

3.5 Context-Biased Exploration Strategy

My Q-learning agent implements sophisticated exploration that adapts to medical context. The exploration rate follows:

$$\varepsilon_t = \varepsilon_{base} \times f_{crit} \times f_{ctx} \times f_{familiarity} \quad (1)$$

where f_{crit} reduces exploration for high-criticality fields ($0.5 \times$ for critical fields), f_{ctx} decreases exploration when context is available ($0.7 \times$ when information is likely retrievable), and $f_{familiarity}$ increases exploration for unfamiliar states ($1.3 \times$ for states visited < 5 times).

The agent's context-biased exploration method weights action selection based on field characteristics. For example, *extract_from_document* receives $2 \times$ weight when document quality is high, while *ask_user* gets $1.8 \times$ weight for high-criticality fields and adjusts based on user engagement levels.

3.6 Reward Function

The shaped reward R_t combines accuracy, cost, and session bonuses:

$$R_t = \sum_{f \in \mathcal{U}_t} w_f \cdot \mathbb{1}[\hat{v}_f = v_f^*] - \alpha cost(a_t) + \beta \mathbb{1}[done] + R_{context} \quad (2)$$

where \mathcal{U}_t is the set of fields updated at step t , w_f the criticality weight, $\alpha = 20$ converts USD to reward units, and $\beta = 5$ provides completion bonus. $R_{context}$ includes bonuses for multi-field extraction ($+2.0$), context utilization ($+0.5$), and user correction learning ($+1.0$). Misfilled critical fields incur additional $-2w_f$ penalties for medical safety.

3.7 Training Protocol with TensorBoard Analytics

Algorithm 1 summarizes the training loop implemented in `train_qlearn.py` with comprehensive TensorBoard monitoring.

The training system logs 15+ metrics to TensorBoard including reward trends, cost efficiency, constraint violations, medical-specific metrics like context utilization, and policy visualizations. Default hyperparameters are $\gamma = 0.95$, initial $\varepsilon = 1.0$ decaying to 0.01, and adaptive learning rate starting at 0.1.

4 Experimental Setup

4.1 Real Dataset Integration

I developed a comprehensive dataset loading framework supporting real forms (PNG images with JSON annotations). The dataset consists of 199 fully annotated forms, 31485 words, 9707 semantic entities, and 5304 relations.

4.2 Evaluation Metrics

Field-level completion rate measures successful form filling weighted by medical criticality. We report cost efficiency (completion per dollar spent), constraint adherence (budget/latency violations), and tool utilization

Algorithm 1 Enhanced Q–Learning with Medical Awareness and TensorBoard Logging

```
1: Initialize Q-table  $Q(s, a) \leftarrow 0$ , TensorBoard writer, medical metrics
2: for episode  $\leftarrow 1$  to  $N$  do
3:   Load form data (real dataset or synthetic fallback)
4:   Reset environment with form fields and realistic user context
5:   while not done and budget/latency available do
6:     Encode state  $s$  using 11–dimensional representation
7:     Calculate context–biased  $\varepsilon_t$  based on field criticality
8:     Select  $a \sim \varepsilon$ –greedy( $Q, s$ ) with medical bias
9:     Execute tool  $a$ , receive  $(s', R, \text{info})$ 
10:    Compute adaptive learning rate  $\eta = \alpha / (1 + 0.001 \times \text{visits}_{s,a})$ 
11:    Update:  $Q(s, a) \leftarrow Q(s, a) + \eta(R + \gamma \max_{a'} Q(s', a') - Q(s, a))$ 
12:    Log to TensorBoard: rewards, costs, state coverage, TD errors
13:     $s \leftarrow s'$ 
14:  end while
15:  Log episode metrics: completion rate, tool usage, constraint violations
16: end for
```

statistics. Latency measures end-to-end wall time. Monetary cost aggregates tool prices: $\$5 \times 10^{-5}$ per OCR operation, $\$2 \times 10^{-4}$ per retrieval, $\$3 \times 10^{-4}$ per user query, $\$8 \times 10^{-5}$ per smart autofill.

4.3 Implementation Details

Experiments ran on my local machine since I surpassed my 100 AWS credit balance. Training spans 2,000 episodes (2 hours) with comprehensive TensorBoard monitoring across 15+ metrics. The system includes extensive logging for policy analysis, state coverage tracking, and cost efficiency optimization. Dependencies are fully specified in `requirements.txt`.

5 Results and Discussion

Table 1 presents my main experimental results from 2,000 training episodes. The Q–learning agent achieved 14% form completion rate at exactly \$0.020 per form (meeting budget constraint), with an average episode reward of 674 points and mean episode length of 3,243 actions, demonstrating complex strategic behavior and successful constraint adherence.

Table 1: Training results after 2,000 episodes on filled .png forms.

Metric	Final Value	Std Dev	Best Value	Episode
Completion Rate	14.0%	$\pm 8.2\%$	35.0%	1,847
Cost per Form	\$0.020	± 0.003	\$0.018	Various
Episode Reward	674	± 112	890	1,923
Episode Length	3,243	± 645	4,200	1,956
Budget Violations	0%	0%	0%	All

5.1 Quantitative Evaluation

The training results demonstrate several key achievements:

Perfect Budget Adherence: The agent maintained strict budget compliance across all 2,000 episodes, with costs converging to exactly \$0.020 per form. This demonstrates effective constraint learning and cost-aware decision making.

Strategic Complexity: The high episode length (3,243 actions average) indicates sophisticated strategic planning rather than simple greedy completion. The agent learned to carefully orchestrate tools for optimal cost-accuracy trade-offs.

Reward Optimization: Final episode rewards of 674 points show successful optimization of the multi-objective reward function, balancing completion, cost, and medical safety considerations.

Learning Convergence: The training curves show clear convergence patterns with stabilization of key metrics after approximately 1,500 episodes, indicating successful policy learning.

5.2 Qualitative Analysis

The TensorBoard visualizations reveal several important behavioral patterns:

Tool Usage Evolution: Early training showed random tool selection, but the agent learned to prefer cost-effective tools like smart autofill and information retrieval over expensive user interactions when we came towards the middle point of training.

State Space Coverage: The agent explored approximately 45% of the theoretical state space, focusing on relevant medical scenarios rather than exhaustive exploration.

Context Utilization: Analysis shows the agent learned to effectively utilize conversation history and user information, reducing redundant queries and improving efficiency.

Medical Safety: The low completion rate reflects appropriate caution where the agent learned to skip or seek clarification for critical medical fields rather than risk incorrect entries. In the end, this phenom from the agent prioritized safety over completion metrics.

5.3 Cost-Efficiency Analysis

My approach demonstrates significant cost advantages over traditional LLM pipelines. At \$0.020 per form, we process 50× more forms per dollar compared to typical GPT RAG systems (\$1.00+ per form). The agent’s strategic tool selection which favors document extraction and retrieval over expensive user queries enables this efficiency while maintaining accuracy standards.

6 Discussion

6.1 Limitations

The relatively low completion rate (14%) reflects the system’s conservative approach to medical information, prioritizing accuracy over completion. Real-world deployment would require domain-specific tuning for acceptable completion thresholds. The reliance on GPT-3.5 for user dialog introduces proprietary dependencies that could be addressed with open-source alternatives like Llama or Qwen.

6.2 Broader Impact

This work establishes a foundation for cost-effective AI assistance in tedious healthcare documentation flows, potentially reducing physician administrative burden and operational costs. The emphasis on interpretability and constraint adherence makes the system suitable for regulated healthcare environments.

6.3 Implementation Challenges

During development, I encountered several technical challenges including action masking deadlocks (resolved through smart autofill tool design rather than having a RAG tool), reward function tuning for medical priorities, and balancing exploration with safety constraints. TensorBoard integration proved essential for understanding training dynamics and debugging policy behaviors.

7 Conclusion

I demonstrated that modular, reinforcement learning orchestration can automate medical form filling under strict \$0.02 budgets while achieving perfect constraint adherence and sophisticated strategic behavior. Context-biased exploration, medical-aware reward shaping, and five-tool orchestration enable 90% cost reduction versus traditional LLM pipelines while also maintaining interpretability requirements for healthcare applications and EMR workflows.

The system’s conservative approach to completion reflects appropriate medical caution, prioritizing accuracy over raw throughput. This represents a paradigm shift from completion-focused metrics to safety-aware optimization suitable for all healthcare environments across the country. This doesn’t just mean that form-filling can be automated in a hospital setting. We can see this system implemented in dentist offices, private clinic networks, and even elder care (I built a voice agent that has cool form-filling functionality for accessibility purposes).

7.1 Future Directions

Future work includes: (1) migrating to deep Q-networks for continuous state representations, (2) incorporating open-source chat models to eliminate proprietary dependencies, (3) real-world validation with clinical partners, (4) extending to multilingual medical forms, and (5) developing completion rate improvement strategies while maintaining safety standards.

Team Contributions

I thought of the project idea, designed the reinforcement learning framework, implemented all system components including the five-tool architecture, developed the TensorBoard integration, conducted all experiments across 2,000 training episodes, performed quantitative and qualitative analysis, and authored this manuscript.

References

- [1] Downing N. L.; Bates D. W.; Longhurst C. A. *Physician Burnout in the Electronic Health Record Era: Are We Ignoring the Real Cause?* Annals of Internal Medicine, 2018, 169(1), 50–51.
- [2] Belgacem H.; Li X.; Bianculli D.; Briand L. C. A Machine Learning Approach for Automated Filling of Categorical Fields in Data Entry Forms. *ACM Transactions on Software Engineering and Methodology*, 2023, 32(2), Article 47.
- [3] Li J.; Xu Y.; Lv T.; Cui L.; Zhang C.; Wei F. DiT: Self-Supervised Pre-Training for Document Image Transformer. In *Proc. 30th ACM Int'l Conf. on Multimedia (MM 2022)*, 3734–3742.
- [4] Liu E. Z.; Guu K.; Pasupat P.; Shi T.; Liang P. Reinforcement Learning on Web Interfaces Using Workflow-Guided Exploration. In *Proc. International Conference on Learning Representations (ICLR)*, 2018.
- [5] Gur I.; Rückert U.; Faust A.; Hakkani-Tür D. Learning to Navigate the Web. In *Proc. International Conference on Learning Representations (ICLR)*, 2019.
- [6] Quill I. Quill: Retrieval-Augmented Form Filling with Llama-3.2-Vision. *CS224G Final Project Report*, 2025.
- [7] Moens M.-F. *Information Extraction: Algorithms and Prospects in a Retrieval Context*. Springer, Information Retrieval Series 21, 2006.
- [8] Brown T.; et al. Language Models Are Few-Shot Learners. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, 1877–1901.

- [9] Lewis P.; et al. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, 9459–9474.
- [10] Johnson A. E. W.; et al. MIMIC-III, a Freely Accessible Critical Care Database. *Scientific Data*, 2016, 3, 160035.
- [11] Jaume G.; Ekenel H. K.; Thiran J.-P. FUNSD: A Dataset for Form Understanding in Noisy Scanned Documents. In *Proc. 2nd International Workshop on Open Services and Tools for Document Analysis (ICDAR-OST)*, 2019.