

Automated Feature Selection and Preprocessing in SageMaker Autopilot for Robust Tabular Modeling

Assignee Research

June 9, 2026

Abstract

This report synthesises findings from 14 peer-reviewed papers addressing the following research question: To what extent do the automatically selected features and preprocessing steps in SageMaker Autopilot's pipelines improve model robustness across different tabular datasets, evaluated using the Amazon. 6 claims were extracted from source literature; 0 were independently verified against retrieved documents. An automated multi-reviewer quality assessment produced a score of 3.3/10. This report is a machine-generated literature synthesis and does not constitute original research.

1 Introduction

This paper examines: Amazon SageMaker Autopilot: a white box AutoML solution at scale. Research question: To what extent do the automatically selected features and preprocessing steps in SageMaker Autopilot's pipelines improve model robustness across different tabular datasets, evaluated using the Amazon Employee Access dataset and comparing against baseline feature engineering methods?.

2 Methodology

Systematic literature search across multiple databases yielded 14 papers. Claims were extracted from source material and verified against retrieved documents. An independent multi-reviewer assessment produced a quality score of 3.3/10.

3 Results

14 papers retrieved. 6 claims extracted; 0 independently verified. Quality review score: 3.3/10.

4 Limitations

This report is a machine-generated literature synthesis and does not constitute original research. Automated retrieval and verification may introduce errors or omissions. Review scores reflect automated assessment, not human peer review. Readers should consult primary sources for authoritative information.

5 Extracted Claims

Claim	Verified	Confidence
The simple ϵ -greedy algorithm works better than other more complicated bandit algorithms, such as EXP3 and Rotting bandi	×	0.08
The system requires 5 finished HP evaluations before using Bayesian Optimization (BO) to ensure the model can be trained	×	0.08
With only 5 random HP evaluations before moving to ϵ -greedy, the algorithm can successfully identify the best pipeline o	×	0.02
For around 80% of datasets, the algorithm's choice is among the top 3 pipelines.	×	0.05
Using learned zero-shot HP configurations instead of random HPs significantly increases the probability of committing to	×	0.04
The meta-model in Amazon SageMaker Autopilot is trained with a collection of datasets and evaluated on a separate test c	×	0.11

References

- <http://arxiv.org/abs/2012.08489v2>
- <http://arxiv.org/abs/2012.08483v2>
- <http://arxiv.org/abs/2510.07569v2>