PAVE: Lazy-MDP based Ensemble to Improve Recall of Product Attribute Extraction Models

Kushal Kumar, Anoop Saladi

1. BACKGROUND :-

- Attribute Extraction (AE) is an important Amazon problem to ensure high catalog quality & power multiple CXs.
- Traditional AE models extract information from product profiles and may suffer from low recall when these are noisy or uninformative (especially *open* attributes).
- These predictions can be improved using product neighbors (see Figure 1).

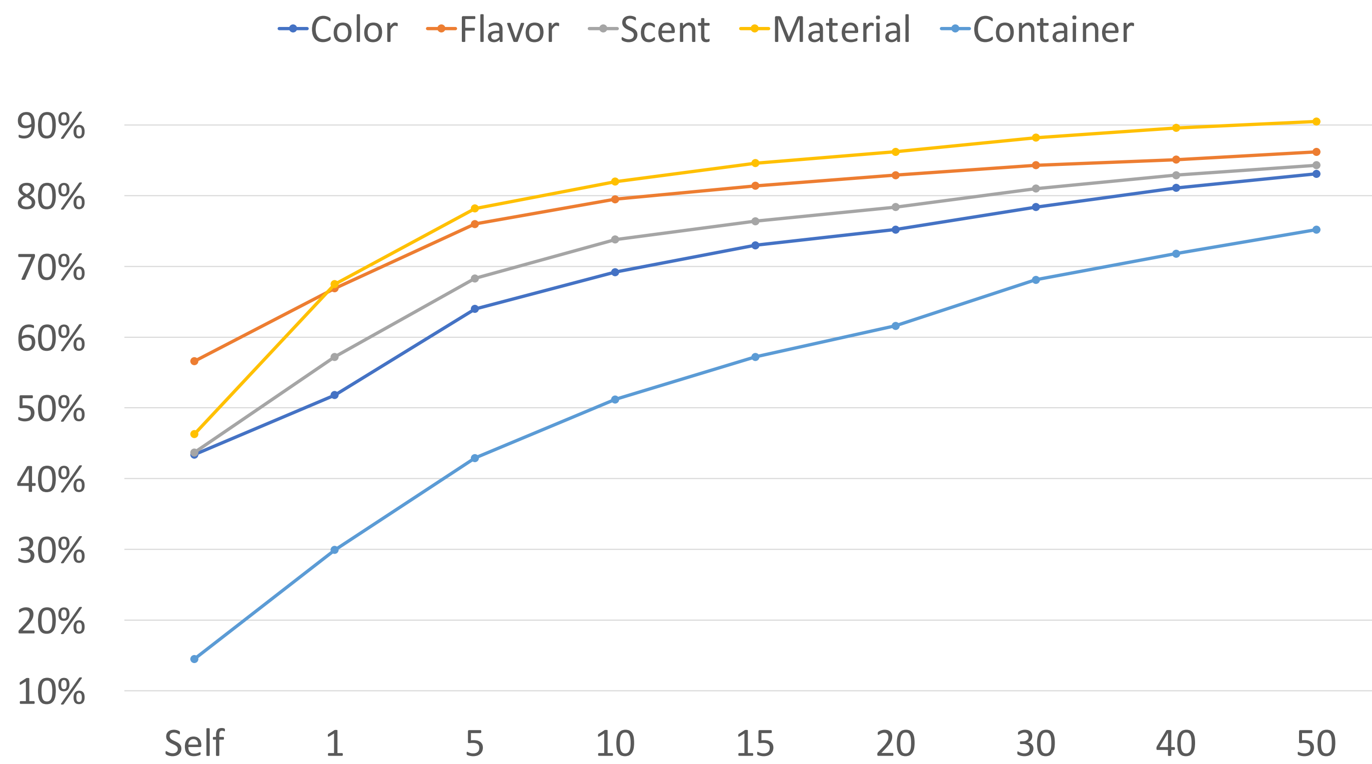

***Figure 1**: Ground truth coverage in neighbor catalog values where product neighbors are ranked from 1-50 basis their cosine similarity of the product embeddings. For Self, we consider both catalog value and baseline AE model prediction. Notably, even the nearest neighbor is only 54% accurate.*

2. PROBLEM FORMULATION :-

- We formulate the problem of attribute value ensemble to choose the correct value from a sequence of product neighbors.
- **Why RL?** – sequential learners have several advantages like handling noise by stopping early, high interpretability, dynamic neighbor length, etc.

3. PAVE :-

- We propose **PAVE: P**roduct **A**ttribute **V**alue **E**nsemble model that uses Lazy-MDP formalism to select the best value from a sequence of product neighbors.
- PAVE lazily scans the sequence of candidate values to ensure that it does not deviate from the AE model often.
- PAVE is a policy network trained using *Proximal Policy Optimization* that was more stable and easy to tune.
- See Figure 3 for training workflow.
- **Default Policy** – simply retains the best value (BV) and skips to next step.
- **Non-lazy actions –** (incurs cost of η)
  - **1** – replace BV with AE Value
  - **2** – replace BV with Vendor Value
  - **3** – replace BV with Current Value
  - **4** – retain BV and end episode
  - **5** – replace BV with blank and end

4. RESULTS :-

- PAVE consistently outperforms BERT-based AE models across 5 attributes* and achieves an average recall lift of 10.3% at the same precision. We apply confidence thresholds to match AE precision using intermediate rewards obtained when the agent selects the final answer.
- See Figure 2 for breakup of recall gain with respect to the neighbor position in the sequence.
- Even on closed attributes like *age range description* and *target gender* where BERT AE models already operate at high recall, PAVE surpasses by 7.8% in precision and 1.5% in recall.
- Non-sequential ensemble techniques like nearest neighbor or majority vote achieve similar recall lift as PAVE but suffer from low precision (17% lower) due to noisy candidates and high bias. PAVE also outperforms model-based ensembleas well.
- PAVE can be scaled by combining multiple attributes & performs well on unseen attributes.

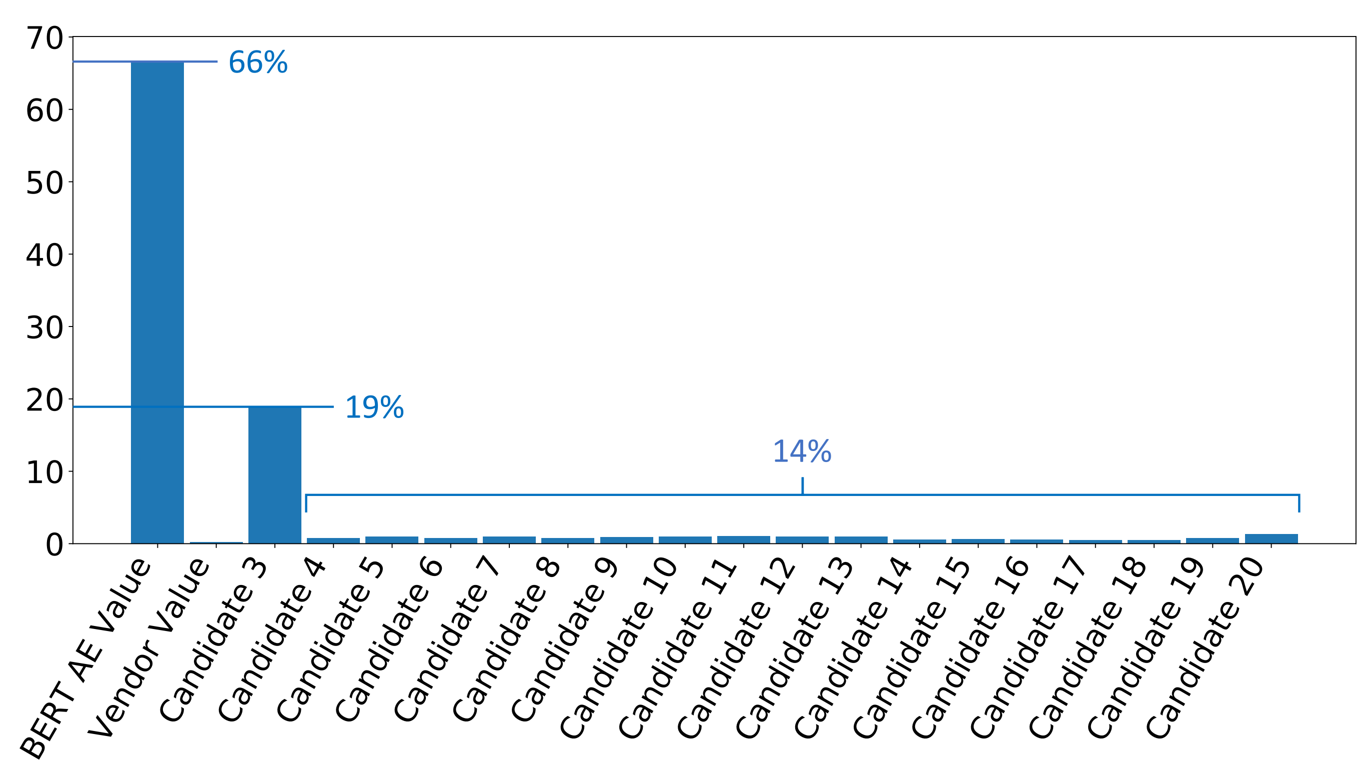

***Figure 2:** Breaking up recall lift from PAVE model by the position of the selected value in the candidate sequence. To plot the graph, we aggregate results for all open attributes. We observe a significant lift coming from neighbors that are uniformly important, hence simple ensemble heuristics fail.*

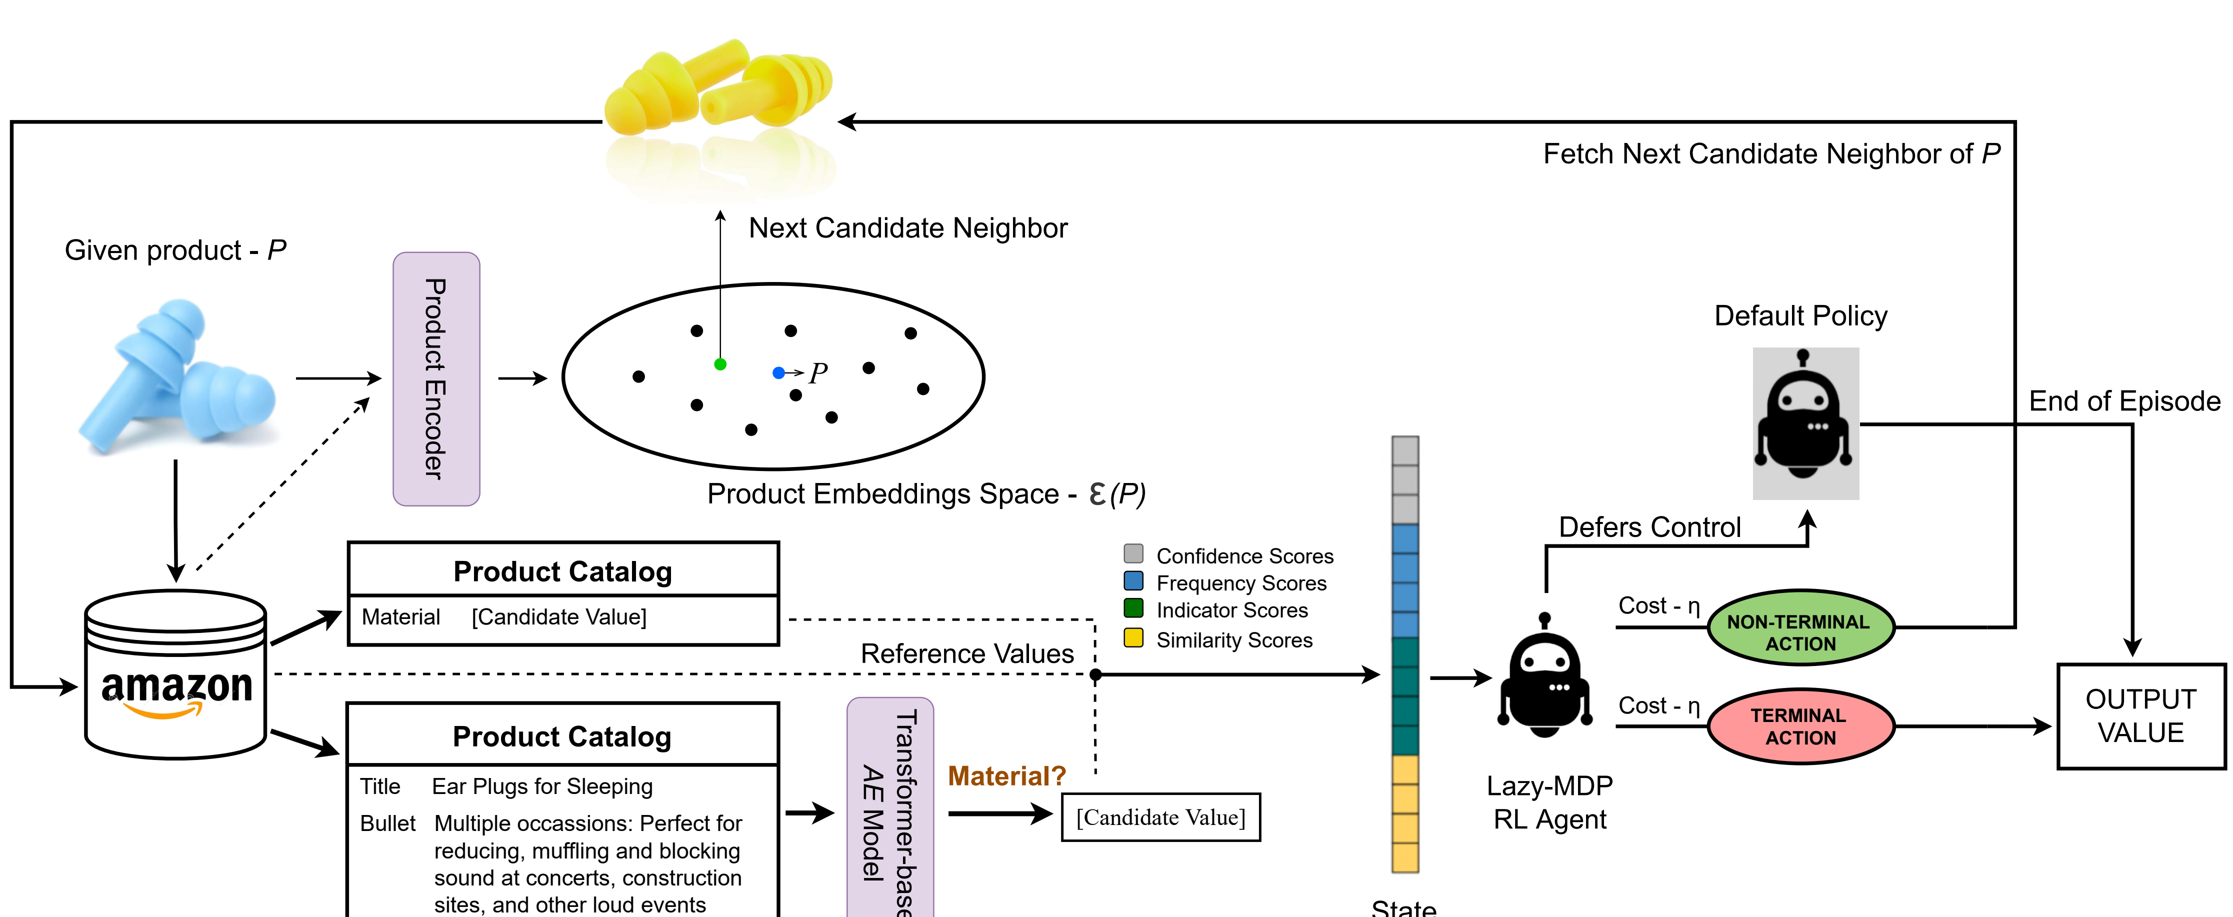

***Figure 3:** Given a product P and attribute say material, catalog value provided by vendor and AE model prediction are used as initial candidate values. PAVE learns to choose the best candidate value by either deferring to a default policy or taking an action at the cost of η. At the end of the episode, the best value is taken as the model output.*

* color, flavor, scent, material and container type