

İpek Sönmez

LOGBOOK

22103939

Week of 2024-10-07 – 2024-10-13

- Began literature review on fake news detection and political bias classification.
- Evaluated possible use cases for CheckMate. Chose to focus on English news only.
- Identified relevant datasets for political inclination and subjectivity.

Week of 2024-10-14 – 2024-10-20

- Investigated the use of LLMs vs custom classifiers.
- Decided on training custom models due to dataset constraints and interpretability needs.
- Outlined bias/objectivity/credibility/similarity pipeline as the architecture for the reliability score.

Week of 2024-11-11 – 2024-11-17

- Adapted an open-source political bias classifier (harshal-11/Bert-political-classification) and modified the head for classification tasks.
- Finalized use case scope and dataset split logic.

Week of 2024-11-18 – 2024-11-24

- Integrated a political bias dataset from Hugging Face/Kaggle into a training pipeline.
- Performed preprocessing, tokenization, and label mapping.

Week of 2024-12-11 – 2024-12-17

- Conducted extensive experimentation with:
- Political inclination classifier (harshal-11/Bert-political-classification)
- Subjectivity classification using a Bi-LSTM with GloVe embeddings
- Semantic text similarity via Google Custom Search + cosine similarity of embeddings
- Evaluated overfitting mitigation strategies including dropout and early stopping.
- Tuned learning rate and batch size for stability (CM-50, CM-51).
- Designed the prototype for the semantic similarity microservice using Google CSE + embedding comparison.

Week of 2024-12-18 – 2024-12-24

- Finalized integration of similarity score thresholds into the reliability model.

- Refactored similarity pipeline to include domain filtering and length validation for candidate articles.
- Benchmarked alternative models (e.g., sentence-transformers/all-distilroberta-v1) for contradiction detection (CM-49).

Week of 2024-12-25 – 2024-12-31

- Assembled a modular reliability scoring pipeline with inputs:
- Bias probabilities (Left, Center, Right)
- Similarity score statistics (mean, max, std, etc.)
- Source credibility (from DB)
- Subjectivity metrics (article and title)
- Spelling error rates
- Planned deployment strategy: EC2 vs. container-based solutions.

Week of 2025-01-01 – 2025-01-07

- Developed and trained a custom MLP-based reliability model using scikit-learn:
- Handled feature scaling with StandardScaler
- Saved pipeline with joblib
- Conducted manual and automated tests using controlled inputs and edge-case samples.

Week of 2025-01-08 – 2025-01-14

- Validated reliability model robustness using:
- Highly subjective, biased, or semantically dissimilar articles.
- Real-world news with varied source credibility.
- Improved scoring stability by thresholding low-confidence similarity scores.

Week of 2025-01-15 – 2025-01-21

- Integrated all three models into a single FastAPI backend:
- /subjectivity: Bi-LSTM model via TensorFlow
- /political: HuggingFace Transformers model (BERT-based)
- /reliability: MLP ensemble with input features
- Wrote detailed documentation for expected input/output formats.

Week of 2025-01-22 – 2025-01-28

- Ran load tests on local deployment and assessed latency per model.
- Created input validation logic for bias\_probs, similarity ranges, and text length.
- Final debugging and reliability regression testing.

Week of 2025-02-01 – 2025-02-07

- Refactored subjectivity classification for both sentence-level and aggregated scoring.
- Optimized memory usage and inference time by parallelizing batch predictions.

Week of 2025-02-08 – 2025-02-14

- Resolved concurrency issues in multi-threaded FastAPI inference with ThreadPoolExecutor.
- Began exploring options beyond EC2 due to AWS Free Tier limitations under concurrent use.

Week of 2025-03-01 – 2025-03-07

- Deployed the ML backend on AWS EC2 using a t2.micro instance.
- Configured instance with gunicorn and uvicorn behind Nginx for routing requests.
- Monitored inference performance under simulated load; observed timeout issues and RAM exhaustion under parallel requests.

Week of 2025-03-08 – 2025-03-14

- Attempted EC2 instance scaling via Elastic Load Balancer (ELB) but was constrained by Free Tier limits.
- Profiled memory consumption and determined model size (especially political classifier) as bottleneck.
- Decided to shift focus to serverless/container-native options.

Week of 2025-03-15 – 2025-03-21

- Containerized inference server pushed to Google Container Registry (GCR).
- Created first Google Cloud Run service deployment with auto-scaling enabled.
- Configured secure environment variables (API keys, scaler paths) using Cloud Run secrets.

Week of 2025-03-22 – 2025-03-31

- Optimized Cold Start time with preloading of models in @startup FastAPI hook.
- Updated FastAPI server to allow CORS requests from frontend and backend.
- Refined /subjectivity, /political, and /reliability endpoints for clearer error handling and logging.
- Confirmed that GCP deployment resolves concurrency and scaling issues—finalized production migration.