Model Monitoring: Detecting and Analyzing Data Issues

Dmitri Melikyan
Founder & CEO at Graphsignal

graphsignal

Model serving observability

- Model serving is a part of production application
- Same observability requirements apply

Data issues: a new failure mode

- Model input schema changes, anomalies and drift
- Input image, text, sound or video specific anomalies
- Missing features, new categories
- Model output anomalies and drift

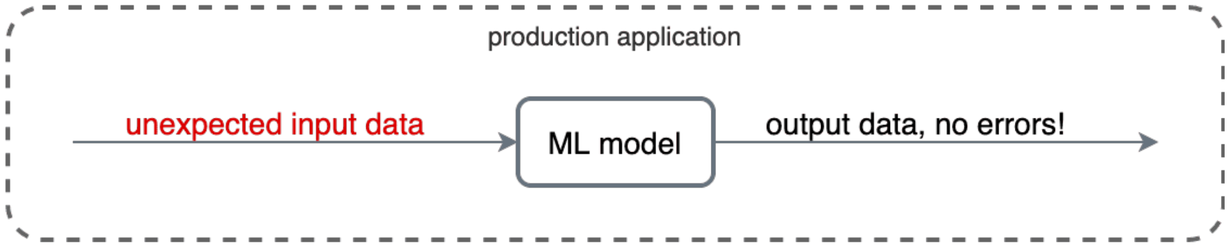

Data dependencies

- Data stores
- Upstream APIs
- Incoming HTTP requests
- Data lakes
- Other sources

Model deployments

- Model deployment is subject to model training process rather than application release cycle
- Model is trained/validated on training data, not on unseen production data

Continuous data monitoring

- Manual, daily or weekly analytics jobs are not optimal for operational issue detection
- Visualization and detection should be optimized for low MTTR

Model monitoring system

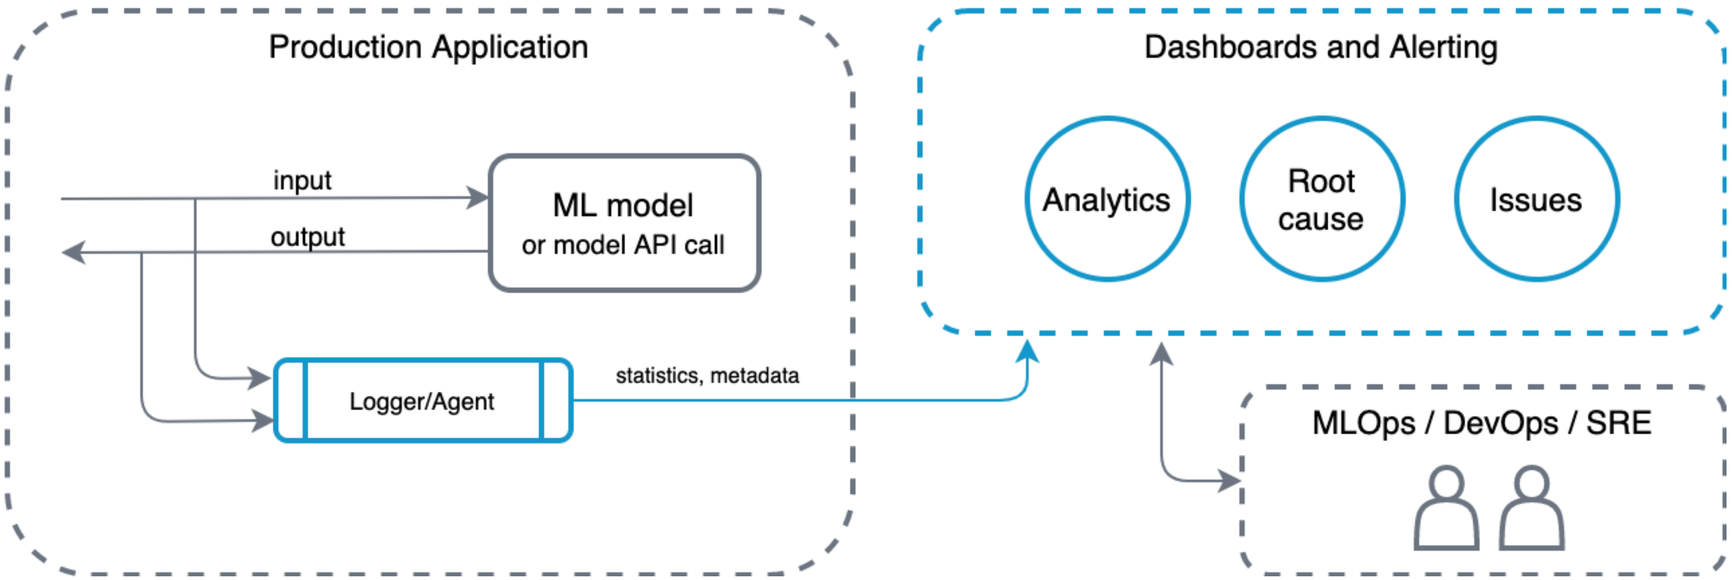

Data analysis dashboard

- Time series for data statistics
- Feature distribution analysis in time

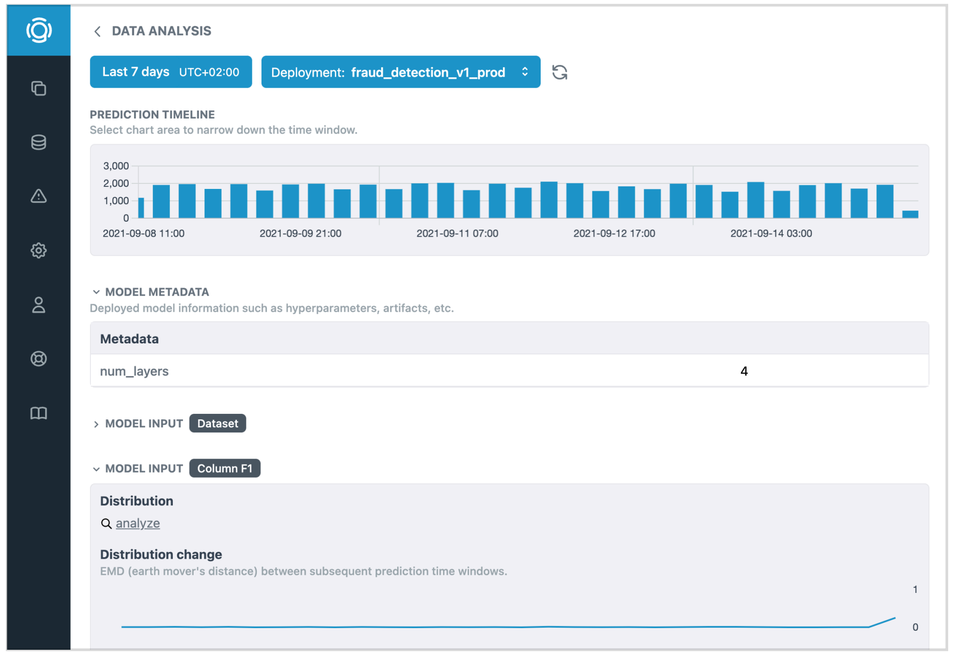

graphsignal

Automatic issue detection

- Too many features to manually setup alerting
- Data specific detection based on computed statistics and baselines

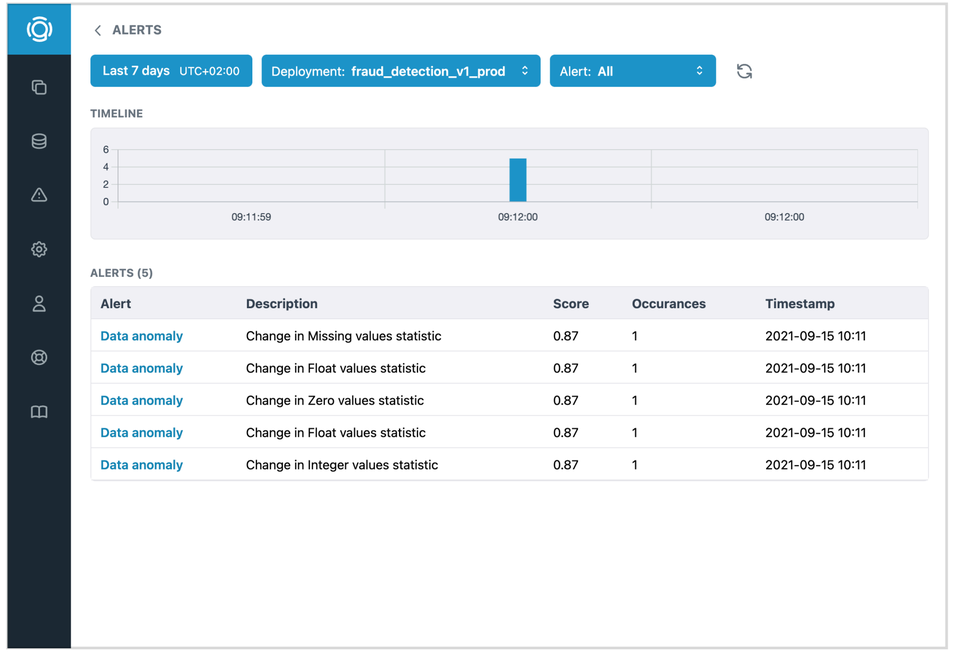

Conclusion

- Data issues in model serving is a new failure mode for ML engineers and MLOps teams to prepare for
- Traditional monitoring systems do not cover the data quality perspective
- Model inputs and outputs should be monitored for validity, anomalies and drift

Thank you!

Questions? dmitri@graphsignal.ai