ML Operations

**Hop maintains ML algorithms in an ongoing manner**

Here at Hop, we build and maintain machine learning infrastructure to enable existing research teams to be more effective in their work.

This often involves setting up compute and storage infrastructure for both raw data and features, as well as systems to track experiments and enable reproducible research. Our operations team members are motivated by efficiency, reproducibility and productivity.

Although bespoke solutions are sometimes necessary, we prefer to assemble them from well-understood (and preferably open-source) parts: Docker, Postgres, Kubernetes/Slurm, Metaflow/Airflow, Databricks, Weights and Biases, etc.

Hop
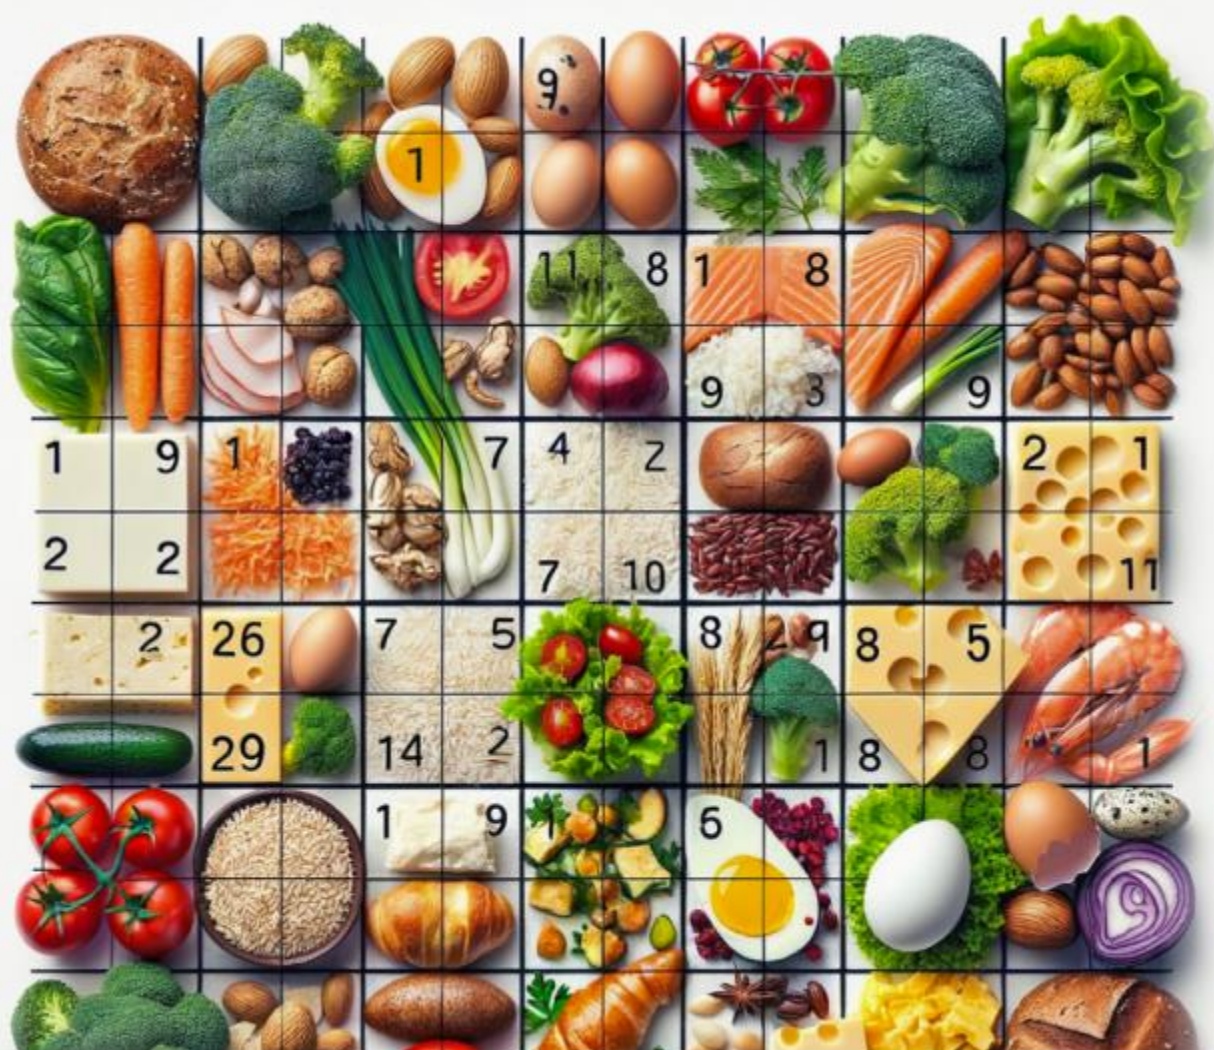
Featured Case Study

Gobble

Improving User Satisfaction via Better Recommendations

As Gobble's trusted partner, we've evolved their recommendation system for a better customer experience, increased sales and reduced churn.

SUMMARY

↘ Gobble offers a number of features in their meal kit delivery service that make them stand out from their competitors.

↘ There are many factors to incorporate for Gobble that are atypical of traditional recommendation systems.

↘ Multiple generations of Gobble's new recommendation system have been rolled out, and work continues to improve the experience for end-users.

↘ Hop was engaged to evolve Gobble's recommendation system as a path toward increasing customer satisfaction and sales, and reducing churn.

↘ Our approach included modernizing the legacy infrastructure and testing a myriad of techniques before moving forward on a system aligned with Gobble's business aims.

READ THE FULL CASE STUDY ↗

Are you looking to leverage ML to improve your operations and business metrics?

Contact us to learn how Hop can help.