



Use case: Predicting and Preventing Dangerous Glucose Events

A prediction task

An Open Health federation could combine signals from the last 24 hours to help predict short-term glucose risks.

Inputs might include:

- **Hospital data:** continuous glucose monitoring (CGM) and insulin use
- **Retailer data:** derived indicators of food purchases and medication adherence
- **Employer data:** derived sleep, activity, and work-shift patterns

The model could then estimate:

- the chance of **severe low blood sugar** in the next 2 hours
- the chance of **significant high blood sugar** in the next 4 hours
- expected **time in range** over the next 24 hours

Longer-term tasks could also be added, such as predicting worsening HbA1c over several months or the risk of complications, but these change more slowly.

What this could look like for the individual

For a person using the system, this might appear as a personalised alert such as:

“Higher risk of overnight low blood sugar tonight based on your recent insulin dose, lower carbohydrate intake than usual, and increased physical activity at work.”

The system could then suggest a timely action, such as:

- having a snack earlier
- reviewing a dose plan if linked to decision support
- receiving a tailored notification at the right moment

Over time, the model could also adapt as the person’s eating habits, work shifts, activity levels, and glucose patterns change.

How the data stay local

The different organisations do not share raw data with each other.

- Each partner trains the model locally on its own systems
- Instead of sending records, they send encrypted model updates



- A central coordinator combines these updates to improve a shared model
- Personalisation layers can then adapt the model to each organisation or individual

This means that health records, shopping histories, and employment data all stay behind each organisation's own firewall, helping meet privacy, regulatory, and commercial requirements.