


Lessons Learnt from Data Exchange Use Cases in the Automotive Network Catena-X

Friday 27th May



- The problem of data privacy
- Our use-case
- What is collective learning? How does it work?
- Differential privacy
- Results



Emma Smith, Collective Learning Team Lead at Fetch.ai

- Achieving state-of-the-art performance on first party data only is unlikely
- Making use of privacy-sensitive datasets is problematic
- Models would be better with this data, which leads to increased business and research value
- Data owners would indirectly benefit from improved models - paid for data access etc.

Shared datasets across company borders provide tremendous value to Machine Learning


Use case “Predictive quality management”

→ Use historical data from multiple assembly/production lines to predict upcoming manufacturing issues

How to train advanced analytics models across multiple companies without centralizing the training data?

Problem statement


- Production data highly sensitive due to intellectual property
- No trustless system to orchestrate ML process
- No incentives



How does it work?

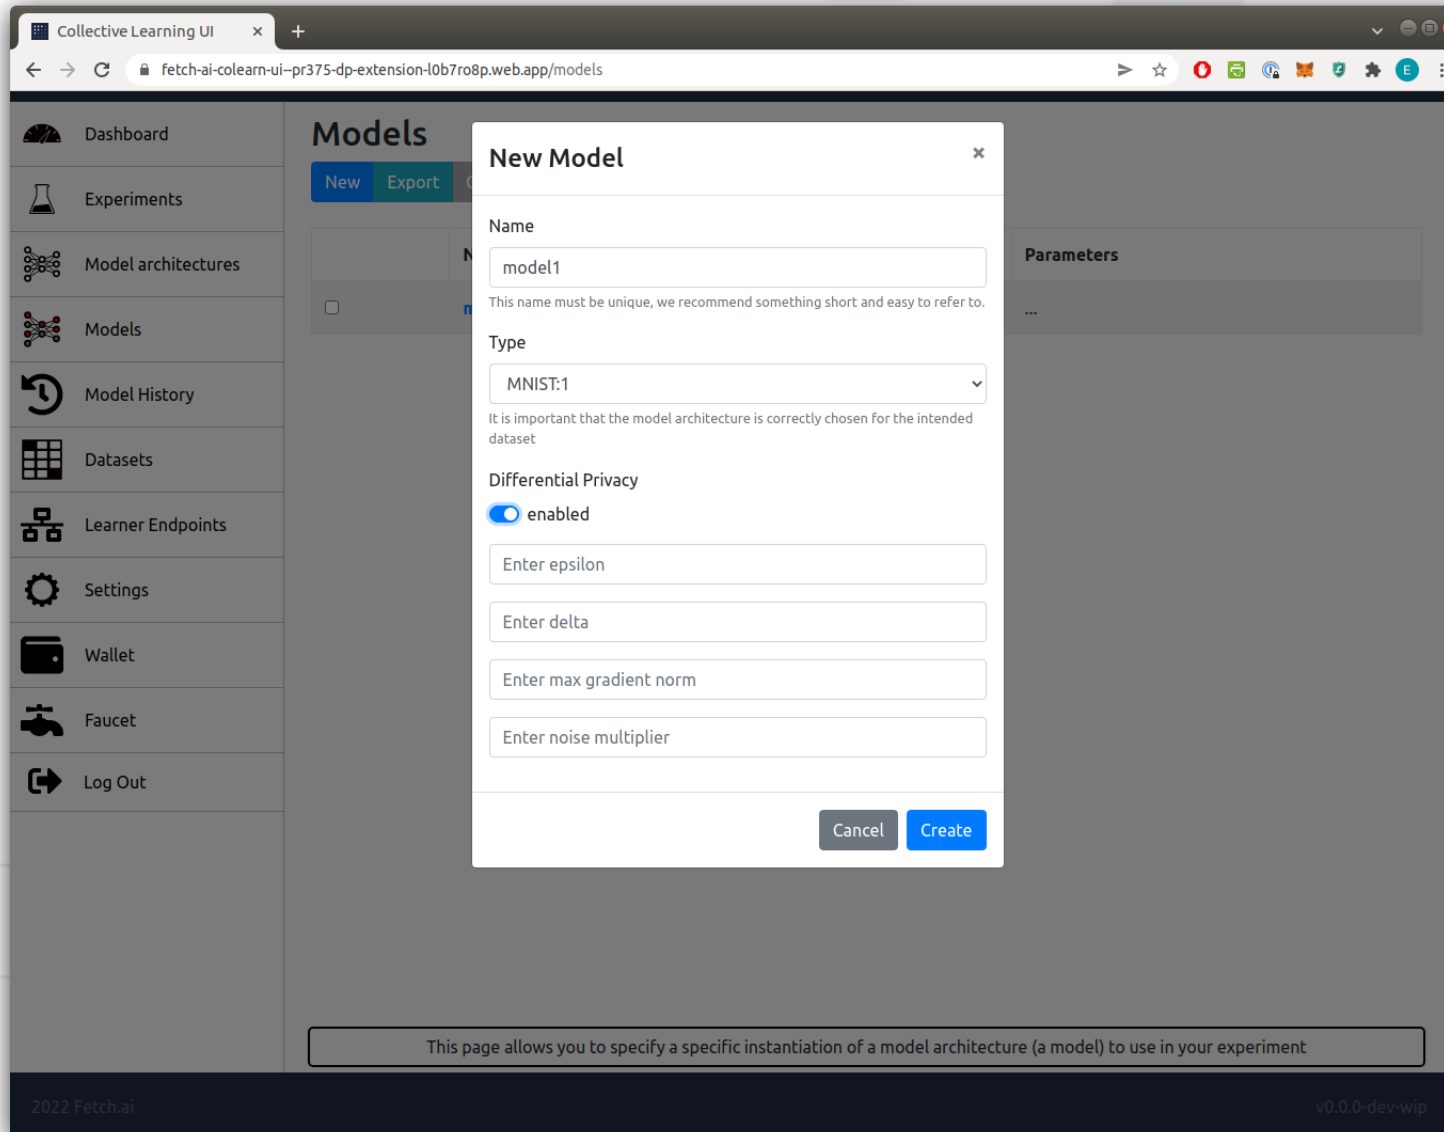
CoLearn Product Suite

 fetch.ai



- Collective learning keeps data local but allows collective training of a model.
- But there is still some leakage - one can work back from a model to reveal information about the training data
- Differential privacy is the solution
- Now enterprises can use private data to train models while retaining control over how much information is revealed

- Differential privacy is a method of limiting the amount of information that is revealed
- Comes from a definition of privacy that focuses on how much an individual's data affects results from the data
- “Privacy budget” - set a limit on how much information can be revealed
- Every time the data is used for training, some of the privacy budget is used up.



The screenshot shows a web browser window titled "Collective Learning UI" with the URL "fetch-ai-colearn-ui-pr375-dp-extension-l0b7ro8p.web.app/models". The main content area is titled "Models" and features a "New Model" dialog box. The dialog box contains the following fields and options:

- Name:** A text input field containing "model1". Below it, a note states: "This name must be unique, we recommend something short and easy to refer to."
- Type:** A dropdown menu currently set to "MNIST:1". Below it, a note states: "It is important that the model architecture is correctly chosen for the intended dataset"
- Differential Privacy:** A toggle switch labeled "enabled" is turned on.
- Input fields for DP parameters:** Four text input fields labeled "Enter epsilon", "Enter delta", "Enter max gradient norm", and "Enter noise multiplier".
- Buttons:** "Cancel" and "Create" buttons at the bottom right of the dialog.

At the bottom of the main interface, a footer bar contains the text: "This page allows you to specify a specific instantiation of a model architecture (a model) to use in your experiment". The footer also includes "2022 Fetch.ai" on the left and "v0.0.0-dev-wip" on the right.



- Colearn system was able to train models and improve accuracy with differential privacy enabled
- Users were able to set and monitor their privacy budgets using a simple UI
- Privacy budget was enforced by system - user left the experiment once budget was used





fetch.ai



Manufacturing - data from production lines can be used to predict when machines need maintenance and predict manufacturing failures

- **Thousands of identical machines**, located **all over the world**
- Machine **manufacturer** is **not allowed** to **access** machine **data**, it belongs to the customer
- **AI data model** is built by the **manufacturer**
- collective models for predictive maintenance or optimal process control
- customers **incentive** to **participate**: No model means high costs when machine breaks
- here blockchain is only visible to machine manufacturer

Connecting data & AI model & Blockchain

- one common AI data model already, but the data remains in separate silos
- one AI collective model can be used by all as a service
- each subsidiary (machine) pays to use the final model
- if the subsidiary (machine) improves the model, it earns credit
- all transactions saved in (private) blockchain
- AI model performance evaluation & KPIs (for businesses)