

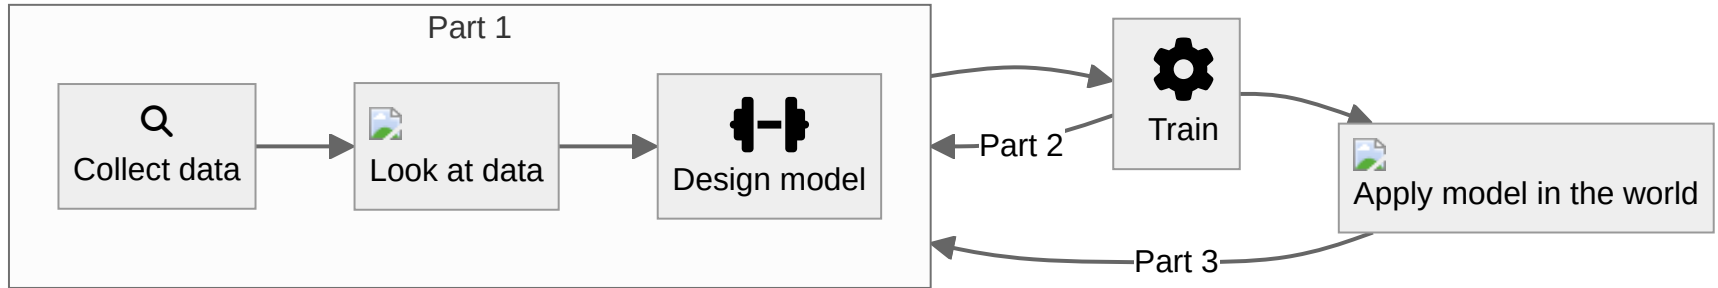
Making It Work - Summary

Making It Work

Part 1: Data and Advanced Network Design

Part 2: Advanced Training

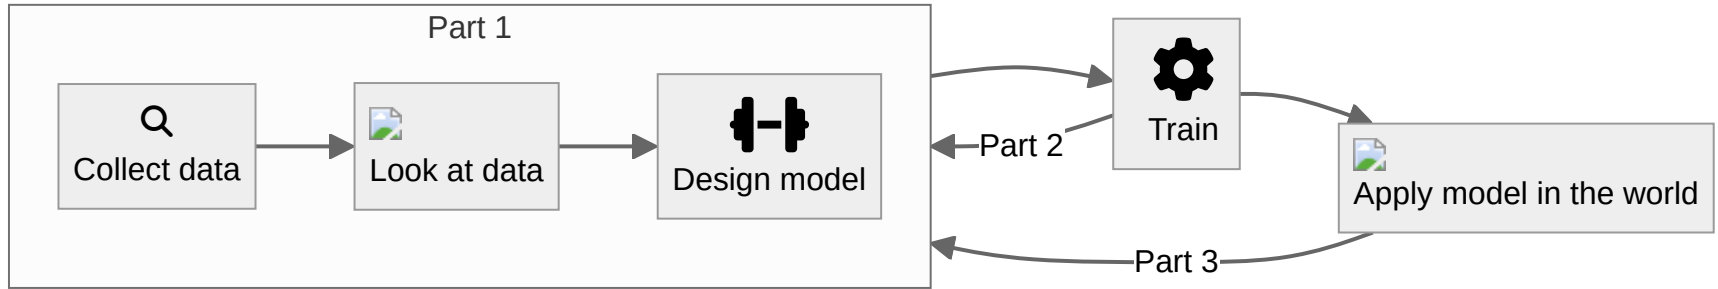
Part 3: Overfitting



Making It Work

Part 1: Data and Advanced Network Design

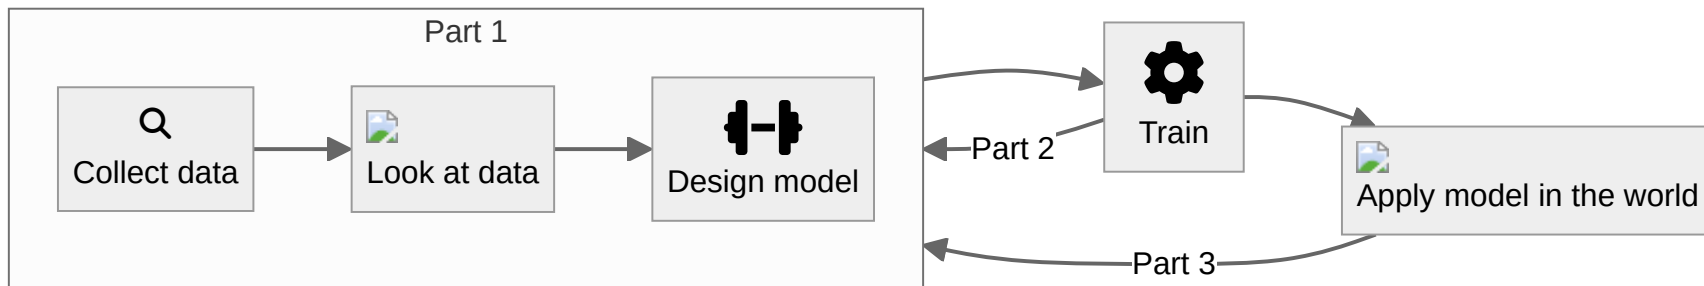
- Look at your data!
- Normalize your inputs



Making It Work

Part 2: Advanced Training

- Adaptively scale per-layer learning rates using AdamW
- Use a learning rate schedule



Making It Work

Part 3: Overfitting

- Create training, validation, and test sets
- Measure overfitting using the validation set
- Prevent overfitting with improved datasets and modeling

