From Rules to Machine Learning to Deep Learning

Laura E. Boucheron

College of Engineering | Klipsch School of Electrical & Computer Engineering

BE BOLD. Shape the Future.

The MNIST Dataset

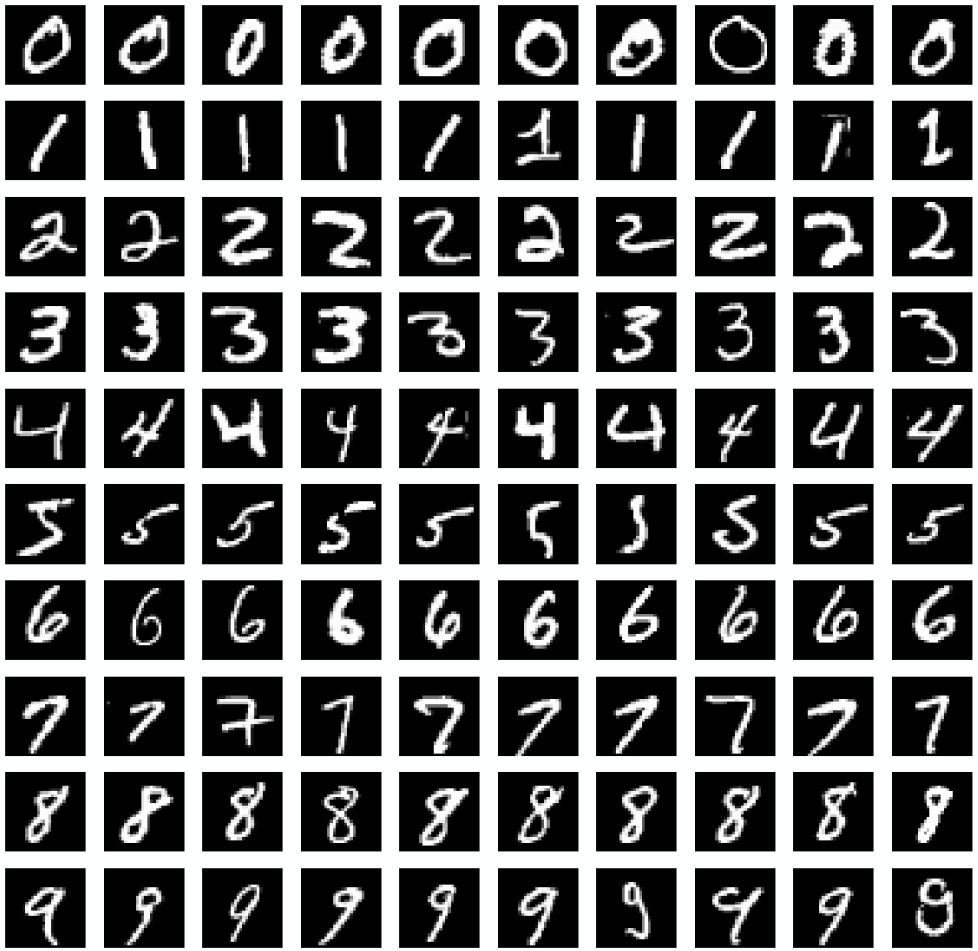

- 70,000 28x28 pixel digitized images of handwritten digits 0 through 9
- Considered a standard benchmark dataset
- Small enough to
  - Fit in memory
  - Run on a modest machine
- Large enough to
  - Span a reasonable range of appearances
  - Solve an interesting problem

Rules:
How might we as humans describe the difference between handwritten digits?

Discriminating b/w 1 and 7

1

7

Discriminating b/w 1 and 7

1

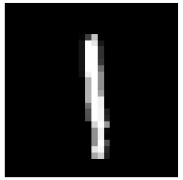

7

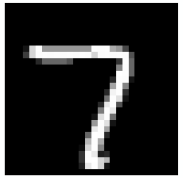

Discriminating b/w 1 and 7

1

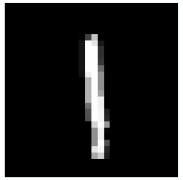

- One line segment

7

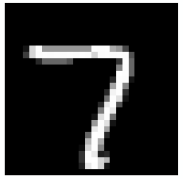

- Two line segments

Discriminating b/w 1 and 7

1

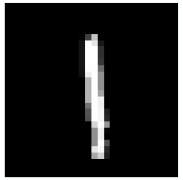

- One line segment

7

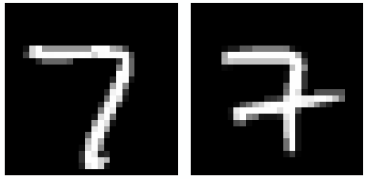

- Two line segments

Discriminating b/w 1 and 7

1

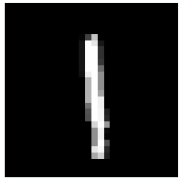

- One line segment

7

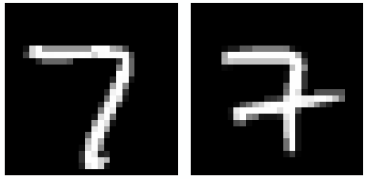

- Two line segments
- OR three line segments

Discriminating b/w 1 and 7

1

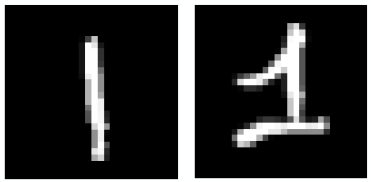

- One line segment

7

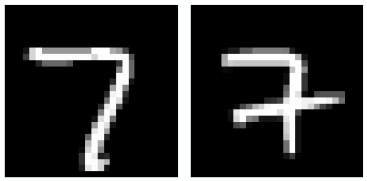

- Two line segments
- OR three line segments

Discriminating b/w 1 and 7

1

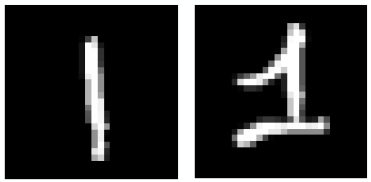

- One line segment
- OR three line segments where two are not perpendicular to the third

7

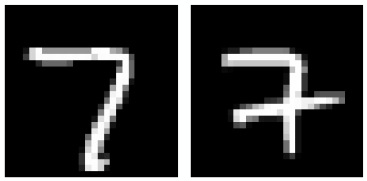

- Two line segments
- OR three line segments

…where two are perpendicular to the third

Discriminating b/w 1 and 7

1

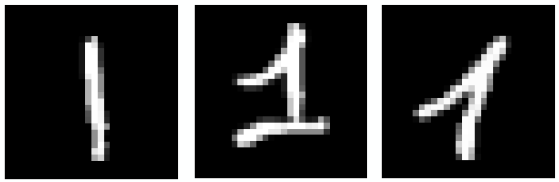

- One line segment
- OR three line segments where two are not perpendicular to the third

7

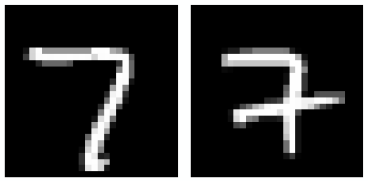

- Two line segments
- OR three line segments

…where two are perpendicular to the third

Discriminating b/w 1 and 7

1

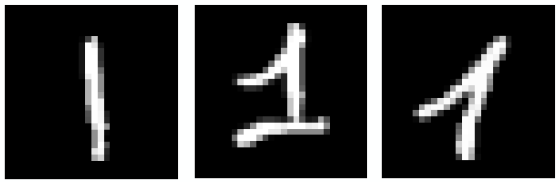

- One line segment
- OR three line segments where two are not perpendicular to the third
- OR two line segments with a small angle between them

7

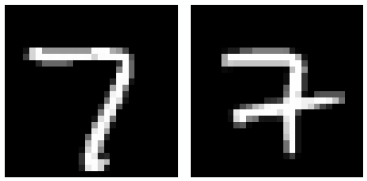

- Two line segments

…with a large angle between them

- OR three line segments

…where two are perpendicular to the third

Discriminating b/w 1 and 7

1

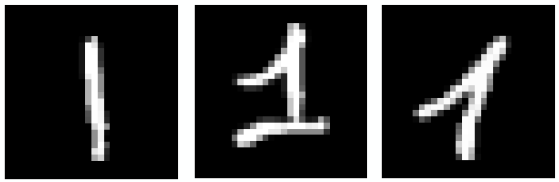

- One line segment
- OR three line segments where two are not perpendicular to the third
- OR two line segments with a small angle between them

7

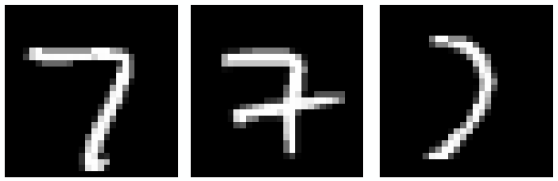

- Two line segments

…with a large angle between them

- OR three line segments

…where two are perpendicular to the third

Discriminating b/w 1 and 7

1

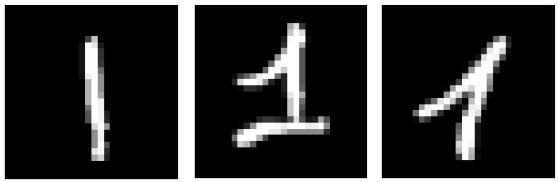

- One line segment

…which has a small curvature to it

- OR three line segments where two are not perpendicular to the third
- OR two line segments with a small angle between them

7

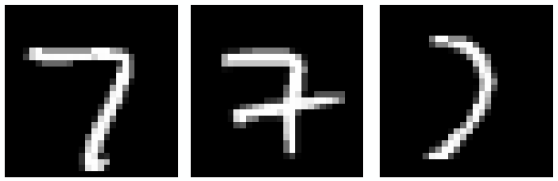

- Two line segments

…with a large angle between them

- OR three line segments

…where two are perpendicular to the third

- OR one line segment which has a large curvature to it

Discriminating b/w 1 and 7

1

- One line segment

…which has a small curvature to it

- OR three line segments where two are not perpendicular to the third
- OR two line segments with a small angle between them

…

7

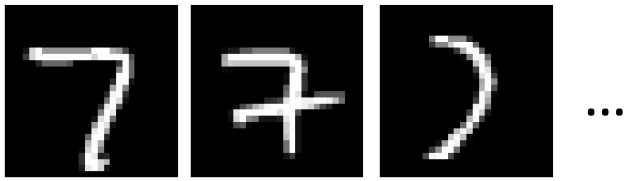

- Two line segments

…with a large angle between them

- OR three line segments

…where two are perpendicular to the third

- OR one line segment which has a large curvature to it

…

Rule-Based Learning

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**

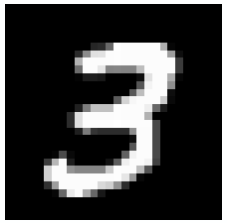

Example Images

"3"

Ground truth

Rule-Based Learning

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **human** to work with specific examples to define a set of discriminatory features—Defines the **feature space**

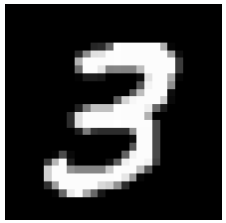

Example Images

"3"

Ground truth

Feature Extraction

- # line segments
- Angle between segments
- Curvature
- # pixels
- Orientation

…

Rule-Based Learning

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **human** to work with specific examples to define a set of discriminatory features—Defines the **feature space**
- Leverage the **human** to use those features to discriminate between digits—Defines the **decision boundary**

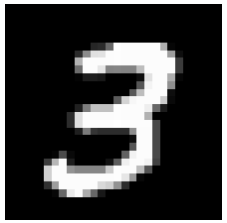

Example Images

"3"

Ground truth

Feature Extraction

- # line segments
- Angle between segments
- Curvature
- # pixels
- Orientation

…

Classification

- Set of rules and thresholds (boundaries)

Rule-Based Learning

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **human** to work with specific examples to define a set of discriminatory features—Defines the **feature space**
- Leverage the **human** to use those features to discriminate between digits—Defines the **decision boundary**

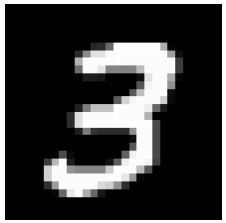

Example Images

"3"

Ground truth

Feature Extraction

- # line segments
- Angle between segments
- Curvature
- # pixels
- Orientation

…

Classification

- Set of rules and thresholds (boundaries)

"3"

Predicted Label

- Leverage **computer** to quickly process many samples

Machine Learning: How can we leverage computers to sift through the features and learn to discriminate between handwritten digits?

Classical Machine Learning: Feature Extraction->Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**

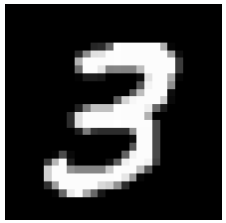

Example Images

"3"

Ground truth

Classical Machine Learning: Feature Extraction->Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **human** to define a set of features that are expected to be discriminatory—Defines the **feature space**

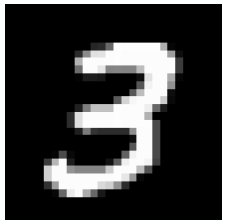

Example Images

Feature Extraction

- # line segments
- Angle between segments
- Curvature
- # pixels
- Orientation

…

"3"

Ground truth

Classical Machine Learning: Feature Extraction->Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **human** to define a set of features that are expected to be discriminatory—Defines the **feature space**
- Leverage the **computer** to learn how to use those features to discriminate between digits—Defines the **decision boundary**

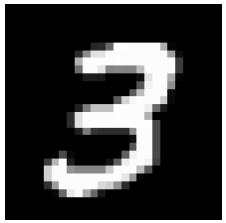

Example Images

Feature Extraction

- # line segments
- Angle between segments
- Curvature
- # pixels
- Orientation

…

Classification

"3"

Ground truth

"3"

Predicted Label

- Leverage **computer** to quickly process many samples

Supervised Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**

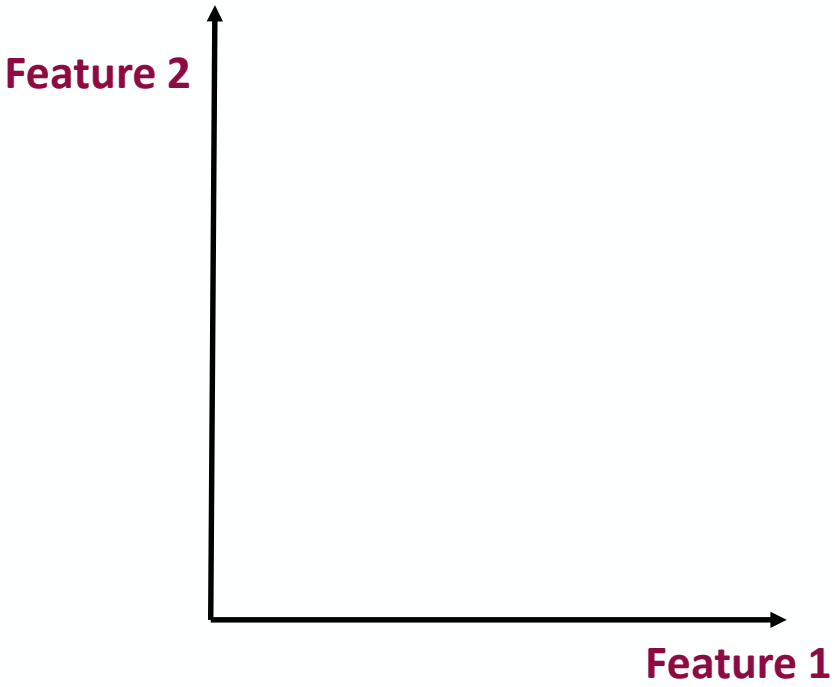

Supervised Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **human** to define a set of features that are expected to be discriminatory—Defines the **feature space**

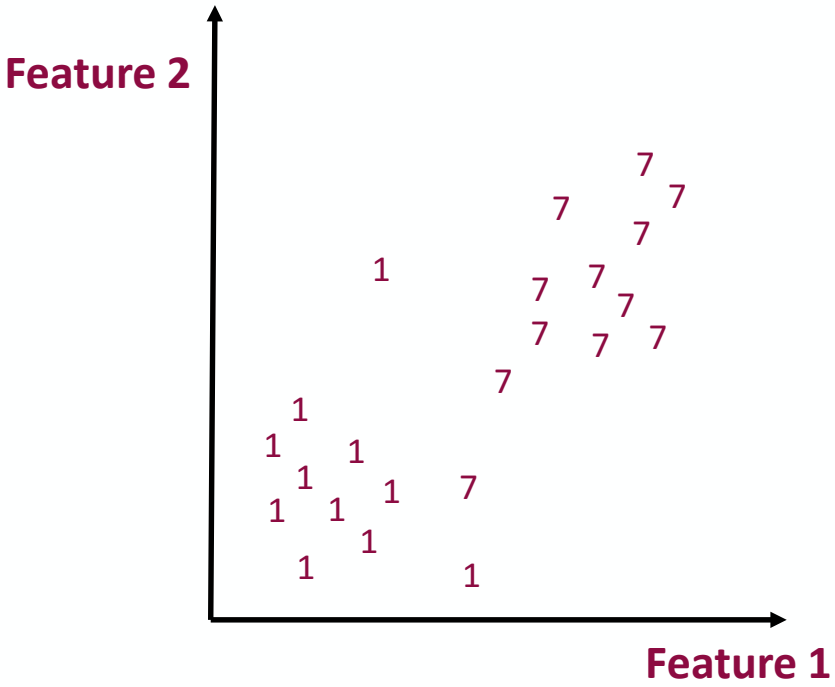

Supervised Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **human** to define a set of features that are expected to be discriminatory—Defines the **feature space**
- Leverage the **computer** to learn how to use those features to discriminate between digits—Defines the **decision boundary**

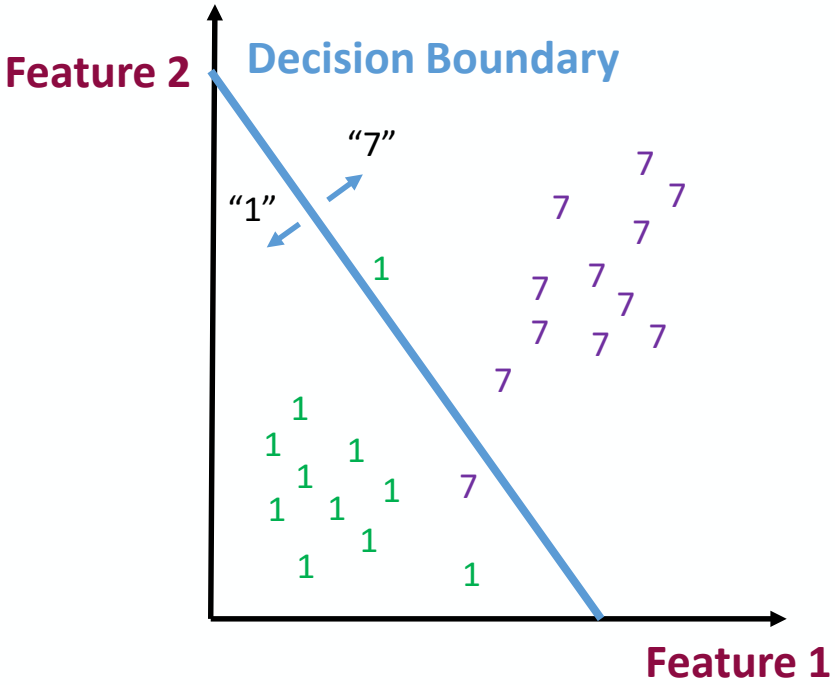

Supervised Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **human** to define a set of features that are expected to be discriminatory—Defines the **feature space**
- Leverage the **computer** to learn how to use those features to discriminate between digits—Defines the **decision boundary**

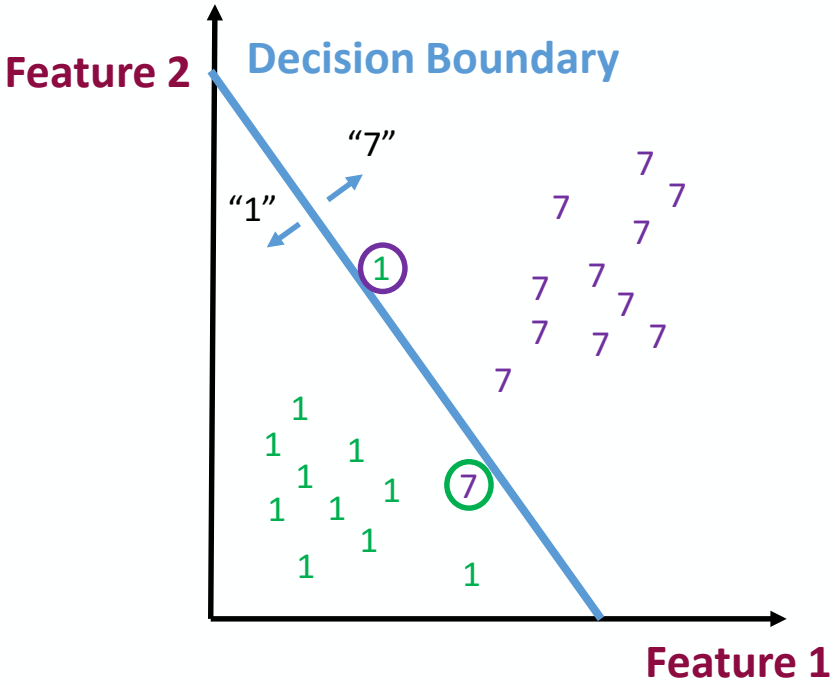

- Misclassifications:
  - Are they due to a **feature space** which is not descriptive enough?
  - Are they due to a **decision boundary** that is not appropriate for the space?
  - Are they due to not enough **training data**?
  - Are they just difficult samples to classify?

Supervised Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **human** to define a set of features that are expected to be discriminatory—Defines the **feature space**
- Leverage the **computer** to learn how to use those features to discriminate between digits—Defines the **decision boundary**

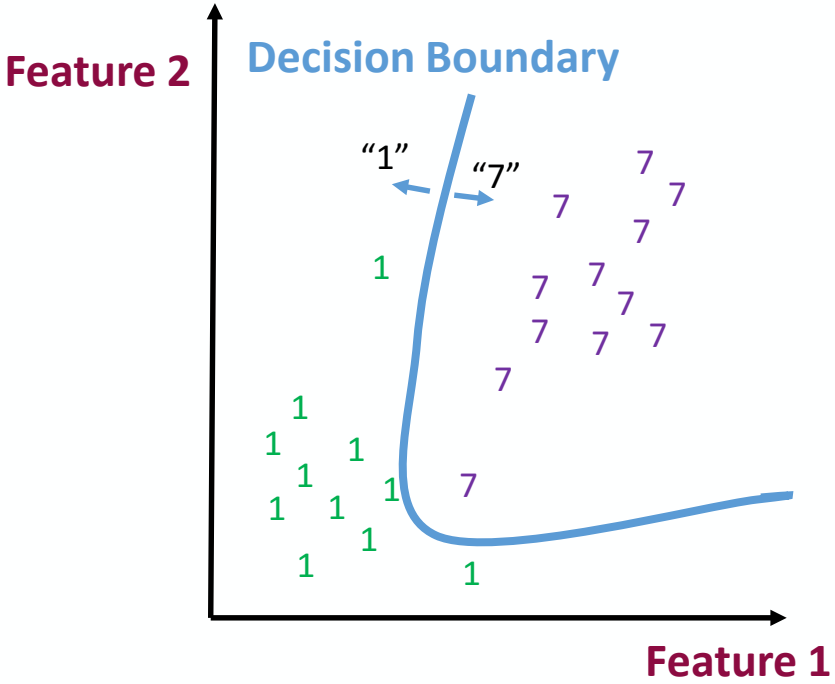

- Misclassifications:
  - Are they due to a **feature space** which is not descriptive enough?
  - Are they due to a **decision boundary** that is not appropriate for the space?
  - Are they due to not enough **training data**?
  - Are they just difficult samples to classify?

Deep Learning: How can we leverage computers to learn the features AND how to use those features to discriminate between handwritten digits?

Deep Learning: Feature Extraction & Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**

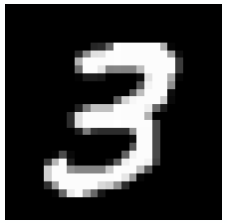

Example Images

"3"

Ground truth

Deep Learning: Feature Extraction & Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **computer** to learn a set of features that are discriminatory—Defines the **feature space**

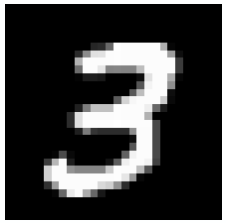

Example Images

"3"

Ground truth

Feature Extraction

Deep Learning: Feature Extraction & Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **computer** to learn a set of features that are discriminatory—Defines the **feature space**
- Leverage the **computer** to learn how to use those features to discriminate between digits—Defines the **decision boundary**

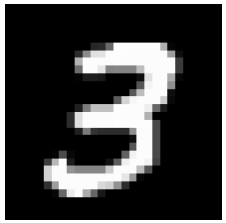

Example Images

"3"

Ground truth

Feature Extraction

Classification

"3"

Predicted Label

- Leverage **computer** to quickly process many samples

Deep Learning: Feature Extraction & Classification

- Leverage the **human** to provide **labeled training data** (example images matched to labels)—Defines the **ground truth**
- Leverage the **computer** to learn a set of features that are discriminatory—Defines the **feature space**
- Leverage the **computer** to learn how to use those features to discriminate between digits—Defines the **decision boundary**

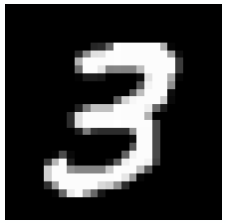

Example Images

"3"

Ground truth

Feature Extraction

Classification

"3"

Predicted Label

Neural Networks…

- Leverage **computer** to quickly process many samples