

Concept Review

29 January 2026

Alex Lyman

Schedule for Today

- Review some of concepts.
- Deep Learning teaser.
- Work a visual example on confusion matrixes and visualizing model error.
- Open Professor hours for last few minutes/leave early.

Principles

- Lots of the things that you're working on might not be something you do in your job (I don't remember the last time I trained a KNN classifier for research) BUT
- The principles of things you learn in this class will be super important. (Model evaluation and validation)
- Later topics (Neural networks, backpropagation, gradient descent, clustering) are all things I use in my life.
- The specific models you will work with might not exist yet!
- Focus on understanding principles.

Model Selection

- How do we know what model to choose?
- One issue is that we only know about 3 types of models. Next week we're doing regression, then trees, then SVMs, then Neural Networks.
- It's hard to fully get into model selection if we only know about 3 models.
- Imagine trying to do construction with a hammer, a screwdriver, and a pair of pliers. (No wrench or saw or power tools)
- Or draw with only 3 crayons!
- Unsatisfying answer – 'You'll understand it better later'

Metrics

- We've looked at a *lot* of metrics.
- Pretty much all of the metrics you need to know can be derived from the confusion matrix.
- Accuracy, Precision, Recall, F1, (Micro and Macro F1), F-Beta, ROC curves, and Precision-Recall curves.
- **All** of these are just fancy ways of counting up confusion matrix stuff.

		POSITIVE	NEGATIVE
		ACTUAL VALUES	
POSITIVE	TP	FN	
NEGATIVE	FP	TN	

$$\text{Precision} = \frac{TP}{TP + FP} \quad \text{Recall} = \frac{TP}{TP + FN}$$

$$\text{Accuracy} = \frac{TP + TN}{TP + FP + FN + TN}$$

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

$$F_{\beta} = (1 + \beta^2) \cdot \frac{\text{Precision} \cdot \text{Recall}}{(\beta^2 \cdot \text{Precision}) + \text{Recall}}$$

$$\text{Macro F1} = \frac{F_1(A) + F_1(B) + F_1(C)}{3}$$

$$\text{Micro Precision} = \frac{\text{Total TP}}{\text{Total TP} + \text{Total FP}}$$


Models

- I want you to have a model-agnostic understanding of Machine Learning.
- A **model** is a mathematical function for describing the relationship between inputs and outputs and making predictions.
- If you can learn (approximate) the function between inputs and outputs, that can be useful in real life.
- We learn this function by tweaking **parameters** inside the model
- How do we know which direction to tweak the parameters? We try to minimize our **error**.

Fake Model Example


How Does a Model Learn?

- A **model** (usually) has trainable *parameters*.
- When training a model, we want to train the parameters using **data**.
- We do this by *minimizing* the error or loss.
- Imagine training a basketball-shooting cannon. You have 2 parameters:
 - Angle (x)
 - Angle (y)




How Does a Model Learn?

- You start with the cannon pointing in some direction.
- How do you know where the cannon will shoot?
- You have to try shooting the cannon.
- You miss. Now what?



How Does a Model Learn?

- You can see how far off you are.
- This is the error (related to loss)
- Then, you make a decision to try and minimize that error.
- In our case, we change our parameters (the x and y angles of the cannon) to try and make a better shot.




How Does a Model Learn?

- We repeat the process:
 - Shoot the cannon
 - Measure how far we missed by (error/loss)
 - Tweak the angle of the cannon (parameters)

until we're making all of our shots

We can think of each time we shoot, measure, tweak as a training example.



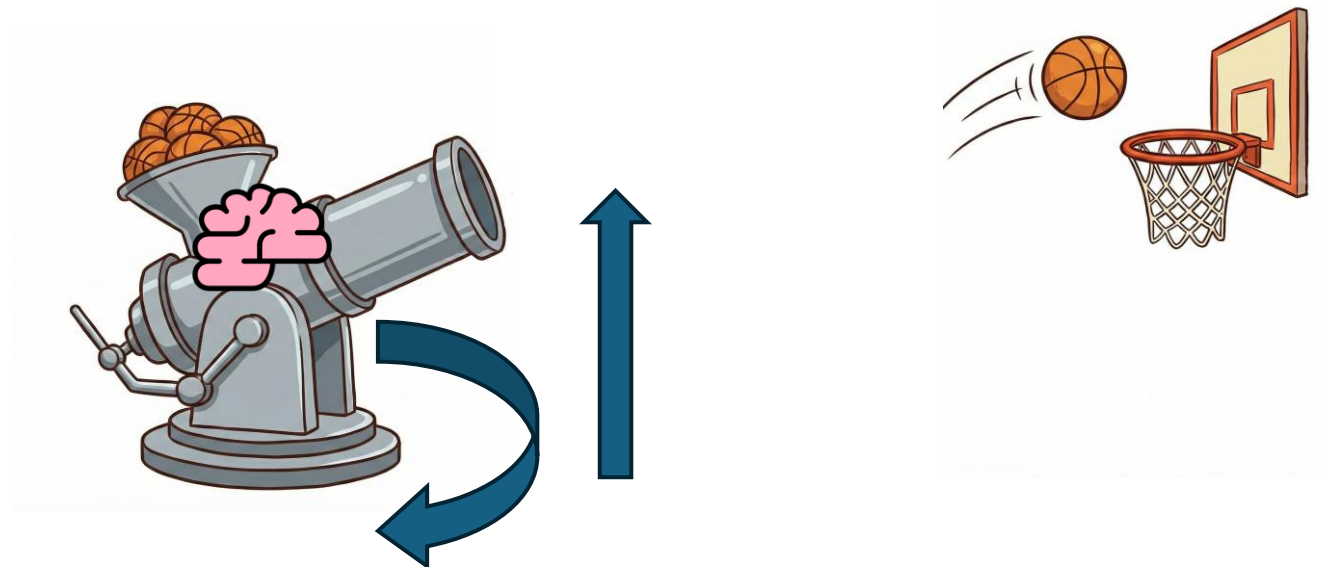
How Does a Model Learn?

- But we don't want to have to tweak the parameters by ourselves.
- Instead, we create some sort of system that tries to minimize error **by itself** by tweaking its own parameters.
- That's the **machine** in **machine learning**.



How Does a Model Learn?

- So instead, we make a little mechanical brain for our cannon. (This is an abstraction). After each shot, we tell it its error, and allow it to tweak its own parameters.
- That's the basic idea of training a machine learning model.




Parameters vs Hyperparameters

Parameters vs Hyperparameters

Feature	Parameters (Weights)	Hyperparameters (Settings)
Who decides?	The Algorithm (The Machine).	The Engineer (You).
When?	During Training (Learning).	Before Training starts.
Analogy	The internal gears of the engine.	The knobs on the dashboard.
Examples	w (Weights), Probabilities.	k (Neighbors), Learning Rate, Smoothing Count.

This Example Again

- Parameters: the X and Y angles of the cannon (chosen by the little mechanical brain)
- Hyperparameter: number of degrees to change at a time



Error/Loss
vs
Accuracy/Precision/Recall/F1

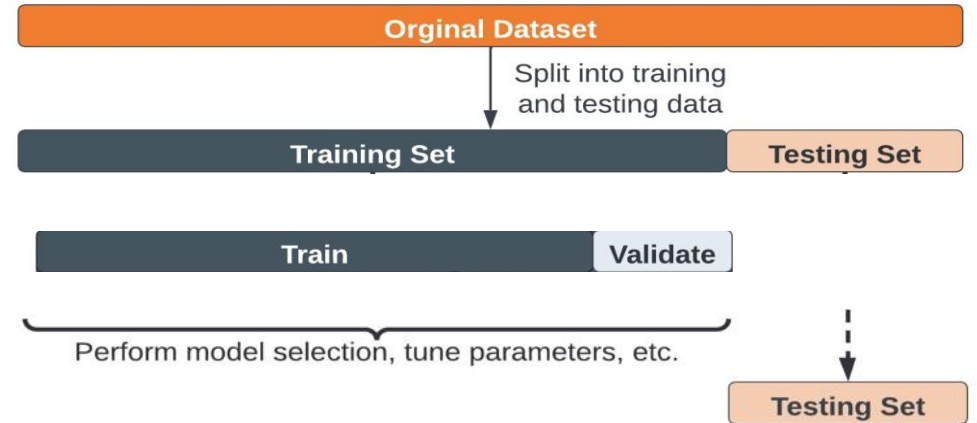
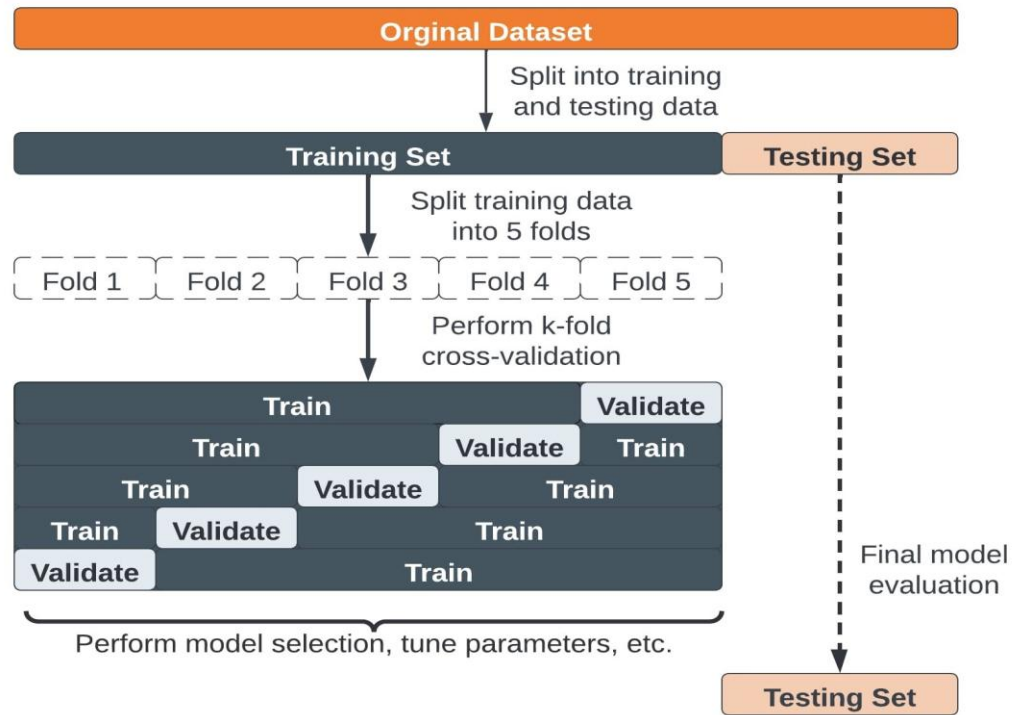
Error/Loss vs Accuracy/Precision/Recall/F1

- We measure the performance of our model multiple times.
- When the model learns, it tries to minimize the difference between its predictions and actual values (the error). It minimizes the error by tweaking its own parameters. This is learning.
- Error and Loss describe the difference between the model's predictions and the targets while the model is training. The *model* cares about error/loss.
- The Confusion Matrix metrics we've talked so much about help us understand how well our model might do in the real world. The *ML Engineer* cares about these metrics.

Training, Validation, Testing

Training, Validation, Testing

- We train a model on data. How can we be sure it will generalize to new data?
- Overfitting (do we remember this?)
- We use an unseen (by the model) Test Set.
- We should ALWAYS use some sort of test set to ensure our model can generalize.
- But how can we make decisions about what hyperparameters to use? How can we decide which model would be best?
- To do that, we use a validation set.
- The validation set can be static, or not (cross-validation), depending on the context.

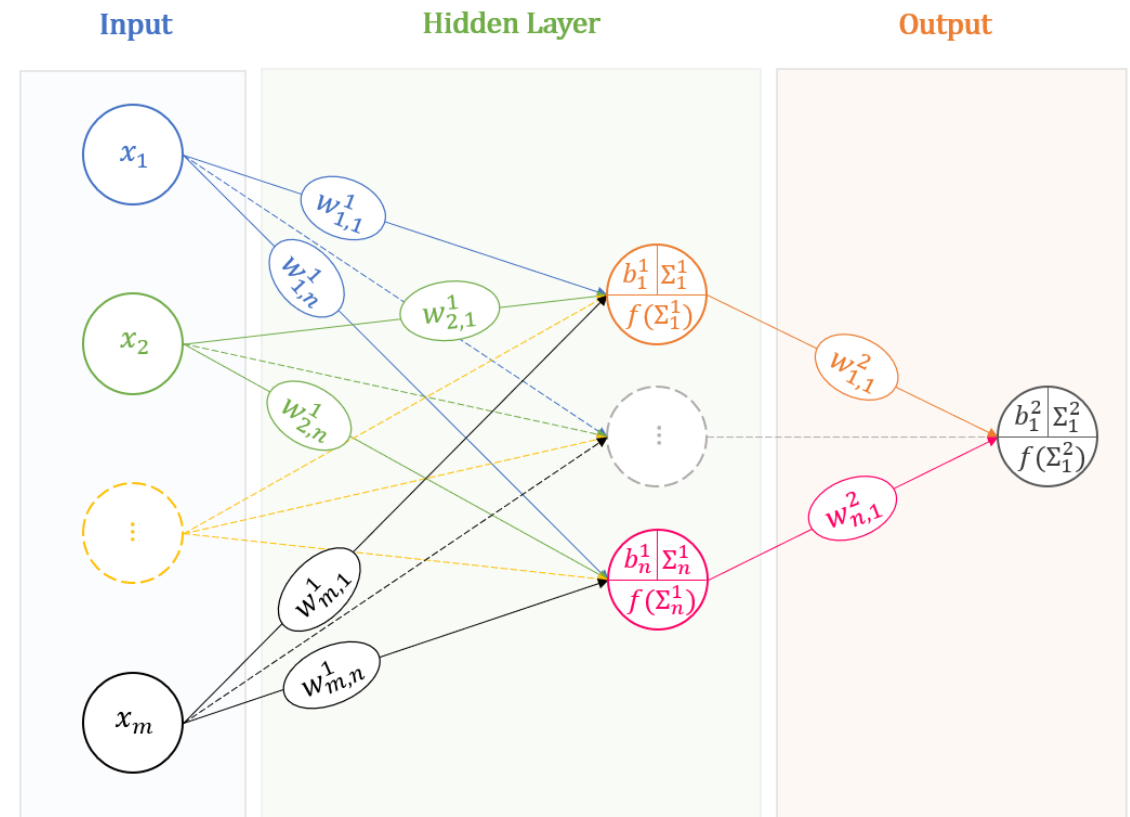


Questions?

Pretraining and Fine Tuning

Deep Learning Basics

- For now, you are going to trust me that an artificial neural net has trainable *parameters* (weights) inside of it.
- You are also going to trust me that the neural net is doing something fancy (gradient descent via backpropagation) to minimize its error/loss.
- Deep learning is **much more computationally intense** than the other methods we will discuss in this class.



Pre-Training and Fine-Tuning

- Imagine if you had two people who had never played soccer before. They've never heard of soccer.
- Your goal is to teach one of them soccer.
- One is a baby and one is an athletic woman in her '20s.
- Who is going to have an easier time learning to play soccer?
- Obviously the woman. She can already run and kick and has coordination.
- The baby would need to learn all of those skills.

Pre-Training and Fine-Tuning

- Training an artificial neural network is a lot more work (computationally) than doing something like logistic regression or KNN.
- When we initialize a neural network, all of the parameters start out randomly initialized.
- The model has to learn the weights by minimizing error over a bunch of training examples.
- That's a lot of work!

Pre-Training and Fine-Tuning

- What if we could do a lot of the work up front one time?
- Pre-training a model involves training the model on generic data so it is capable, before fine-tuning for a specific task.

Pre-Training and Fine-Tuning Example

- Example: Middle English and Modern English are pretty different.
- (word order in line 70, vocab line 75)
- I wanted to make a translator to turn Modern English into Middle English.
- How could I do that?
- There's not enough training data!

70 **He nevere yet no vileynye ne sayde**
He never yet said any rude word

71 **In al his lyf unto no maner wight.**
In all his life unto any sort of person.

72 **He was a verray, parfit gentil knyght.**
He was a truly perfect, noble knight.

73 **But for to tellen yow of his array,**
But to tell you of his clothing,

74 **His hors were goode, but he was nat gay.**
His horses were good, but he was not gaily dressed.

75 **Of fustian he wered a gypon**
He wore a tunic of coarse cloth

76 **Al bismotered with his habergeon,**
All stained (with rust) by his coat of mail,

77 **For he was late ycome from his viage,**
For he was recently come (back) from his expedition,

78 **And wente for to doon his pilgrymage.**
And went to do his pilgrimage.

Pre-Training and Fine-Tuning

- We start with BART, a language model from Meta was trained to learn English from 160 GB of text including:
 - **BooksCorpus:** A large dataset of unpublished books.
 - **English Wikipedia:** The full English Wikipedia (excluding talk pages).
 - **News Data:** CC-NEWS, which is a collection of news articles.
 - **Web Text:** OpenWebText.
 - **Stories:** A dataset of stories
- The pre-training took 16 NVIDIA V100 GPUs (back in 2019 that was a lot).
- This would have cost them thousands of dollars and taken several days of continuous training.
- Then they made it available to the public for free!

Pre-Training and Fine-Tuning

- I gathered about 60,000 paired sentences (Modern and Middle English), taken from Chaucer's complete works, the Wycliffe Bible, and Sir Gwain and the Green Knight.
- Then, I fine-tuned the model on the paired sentences (one gpu, a few hours)

```
{ "en": "And God said, Let the earth put forth grass, herbs yielding seed, `and' fruit-trees bearing fruit after their kind, wherein is the seed thereof, upon the earth: and it was so.", "me": "and seide, The erthe brynge forth..
```

```
{ "en": "And the earth brought forth grass, herbs yielding seed after their kind, and trees bearing fruit, wherein is the seed thereof, after their kind: and God saw that it was good.", "me": "And the erthe brouyte forth greene..
```

```
{ "en": "And there was evening and there was morning, a third day.", "me": "And the euentid and morwetid was maad, the thridde dai." }
```

```
{ "en": "And God said, Let there be lights in the firmament of heaven to divide the day from the night; and let them be for signs, and for seasons, and for days and years:", "me": "Forsothe God seide, Liytis be maad in the firmament..
```

```
{ "en": "and let them be for lights in the firmament of heaven to give light upon the earth: and it was so.", "me": "and shyne tho in the firmament of heuene, and liytne tho the erthe; and it was doon so." }
```

```
{ "en": "And God made the two great lights; the greater light to rule the day, and the lesser light to rule the night: `he made' the stars also.", "me": "And God made twei grete liytis, the gretter liyt that it schulde be bifor..
```

```
{ "en": "And God set them in the firmament of heaven to give light upon the earth,", "me": "and God made sterris; and settide tho in the firmament of heuene, that tho schulden schyne on erthe," }
```

Pre-Training and Fine-Tuning

English to Middle English Translator

This translator is trained on about 70,000 English/Middle English paired sentences.
It's still a work in progress.

sentence

If it had not been for Cotton-Eye Joe, I'd been married long time ago
Where did you come from, where did you go?
Where did you come from, Cotton-Eye Joe?

Clear Submit

output

If it nere for Cotton-Eyen Joe, I sholde been wedded yoore
Where camest thou fro, where wentest thou?
Where camest thou fro, Cotton-Eyen Joe?

Modern English to Middle English Translator

This translator is trained on about 70,000 English/Middle English paired sentences.
It's still a work in progress.

sentence

The sun did not shine
It was too wet to play
So we sat in the house
All that cold, cold, wet day

Clear Submit

output

The sonne ne shoon
Hit was to wete for to pleye
So seten we in the hous,
Al that colde, coold, wete day

Break Time

My model on training data




My model on test dataset



Visualizing a Confusion Matrix With Image Classification

Magic: the Gathering



Magic: the Gathering



- **White** represents order and community, building large armies of small creatures and using protective magic to control the battlefield and enforce its laws.
- **Blue** embodies intellect and the pursuit of perfection, controlling the game by drawing extra cards to gain knowledge and canceling an opponent's spells before they can even take effect.
- **Black** is the color of ambition and amorality, achieving power by any means necessary, whether it's destroying opposing creatures, reanimating the dead, or sacrificing its own life for a winning advantage.
- **Red** channels passion and chaos to act on pure impulse, overwhelming opponents with fast, aggressive creatures and spells that deal direct damage to any target.
- **Green** is the color of nature and growth, embracing the law of the jungle by accelerating its resources to summon gigantic, powerful creatures that overpower the opposition.

Train Time!

White



Blue



Black



Red



Green



Test Time!

Test #1



Test #2




Test #3



Start with a Pretrained Model

- Resnet-18 (special type of neural network specifically for image classification)
- Pre-trained on ImageNet-1k (Over 1 million examples from 1,000 classes)



```
# Setup Model  
model = models.resnet18(weights=models.ResNet18_Weights.IMAGENET1K_V1)
```

Create a Dataset

- Scryfall API
- Get 500 random cards in each color
- 80-20 train/test split (we should get about 100 cards from each color in the test set).

```
# 1. DOWNLOADER FUNCTION
def download_dataset():
    if os.path.exists(DATA_DIR):
        print(f"Directory {DATA_DIR} exists. Skipping download.")
        return

    os.makedirs(DATA_DIR, exist_ok=True)

    for color in CLASSES:
        # Create class subfolder (e.g., ./mtg_data/W)
        class_dir = os.path.join(DATA_DIR, color)
        os.makedirs(class_dir, exist_ok=True)

        print(f"Fetching {color} cards...")

        # Scryfall query: Single color, not digital, has art
        query = f"c={color} -c:m is:spell -is:digital"
        api_url = "https://api.scryfall.com/cards/search"

        count = 0
        page = 1

        while count < SAMPLES_PER_CLASS:
            params = {'q': query, 'page': page, 'format': 'json'}
            resp = requests.get(api_url, params=params)

            if resp.status_code != 200:
                break

            data = resp.json()
            if 'data' not in data: break

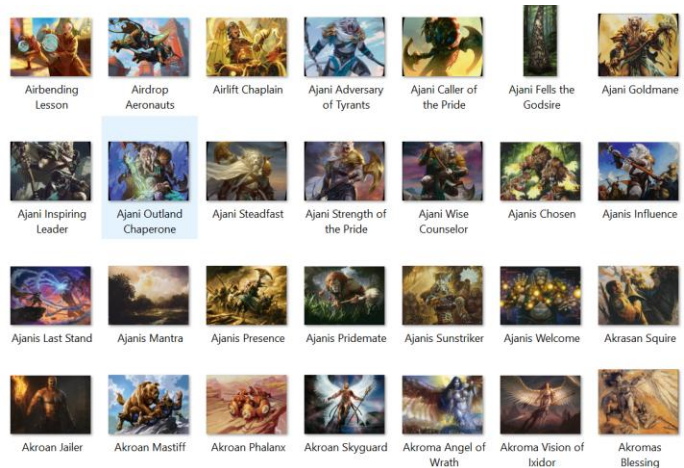
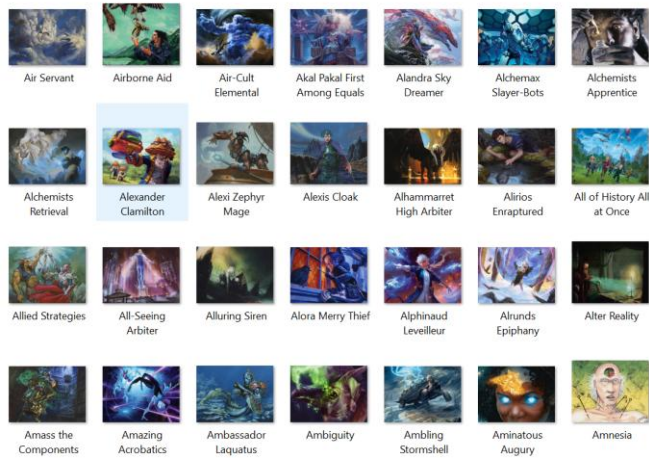
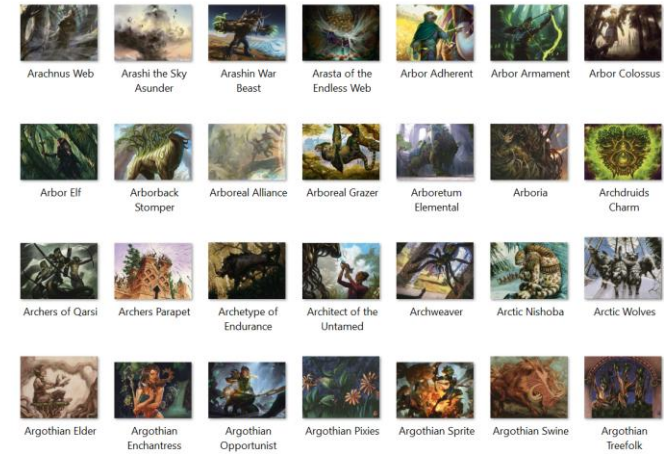
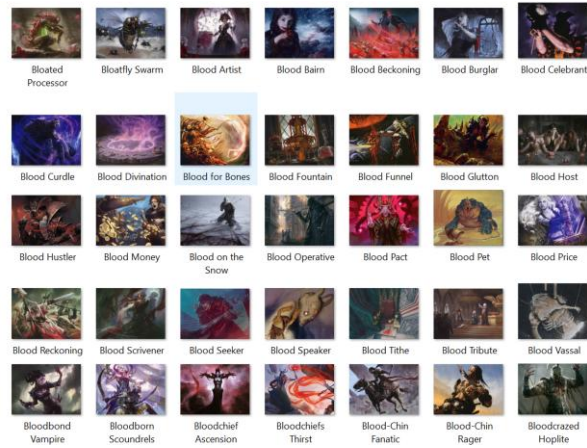
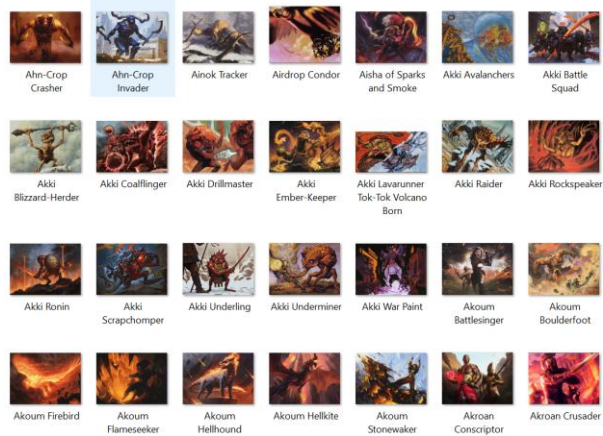
            for card in data['data']:
                if count >= SAMPLES_PER_CLASS: break

                # Safety check for image dict
                if 'image_uris' in card and 'art_crop' in card['image_uris']:
                    img_url = card['image_uris']['art_crop']
                    # Sanitize filename
                    safe_name = "".join([c for c in card['name'] if c.isalnum() or c in (' ', '-')]).strip()
                    save_path = os.path.join(class_dir, f"{safe_name}.jpg")

                    try:
                        img_data = requests.get(img_url).content
                        with open(save_path, 'wb') as f:
                            f.write(img_data)
                        count += 1
                    except Exception as e:
                        print(f"Failed {safe_name}: {e}")

            page += 1
            time.sleep(0.1) # Respect API limits
```

Look at the Dataset



Train the Model

- We train the model on the train set.
- As the model trains, it tweaks its own weights to learn to classify Magic card art based on color.
- Hey look! We're using Cross Entropy Loss! We've heard of that.

```
criterion = nn.CrossEntropyLoss()
optimizer = optim.Adam(model.fc.parameters(), lr=0.001)

print("Training model...")
for epoch in range(5): # 5 epochs is plenty for this
    model.train()
    running_loss = 0.0
    for inputs, labels in tqdm(loader, desc=f"Epoch {epoch+1}"):
        inputs, labels = inputs.to(device), labels.to(device)
        optimizer.zero_grad()
        outputs = model(inputs)
        loss = criterion(outputs, labels)
        loss.backward()
        optimizer.step()
        running_loss += loss.item()

# Save Model
torch.save(model.state_dict(), MODEL_PATH)
print(f"Model saved to {MODEL_PATH}")
```

Test the Model

- We test the model on the test set.
- How well did we do?
- What is random chance?
- How balanced are our classes?
- How well did we do on each class?


```
Classification Report:
              precision    recall  f1-score   support

     B         0.47         0.58         0.52         92
     G         0.60         0.77         0.67        108
     R         0.62         0.43         0.51         91
     U         0.63         0.38         0.47        110
     W         0.50         0.61         0.55         99

 macro avg         0.56         0.55         0.54        500
weighted avg         0.57         0.55         0.55        500
```

Test the Model – Confusion matrix

- Which color was our model best at?
- Worst?
- What kind of errors did the model make?



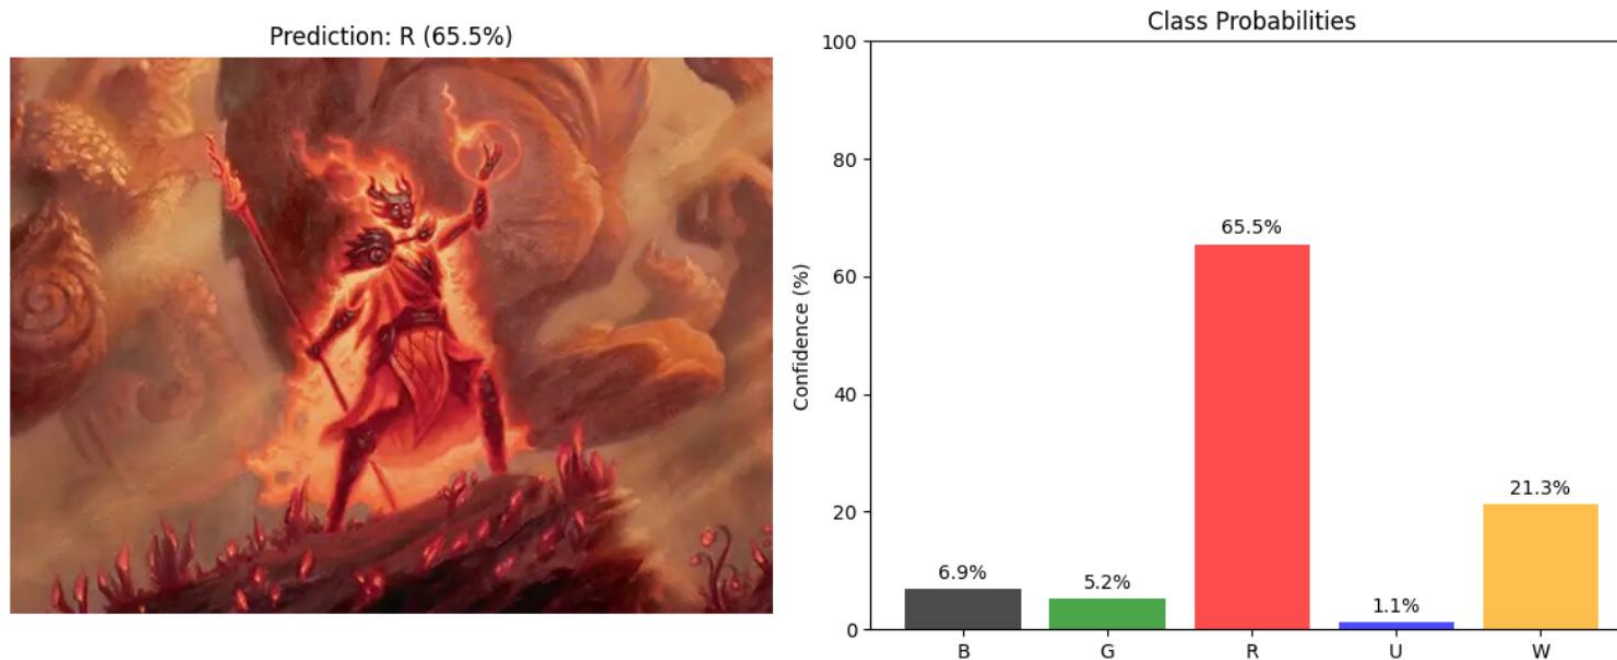
Test the Model – Look at Good/Bad Classifications

- Models can do poorly for a couple of reasons:
 - The model didn't learn what it was supposed to.
 - The task it was learning was messy. (Noise/Variance)
- If our model is failing, but we can be more confident that it is failing for the second reason, that's good.

Let's look at some examples of good and bad classifications by our model:

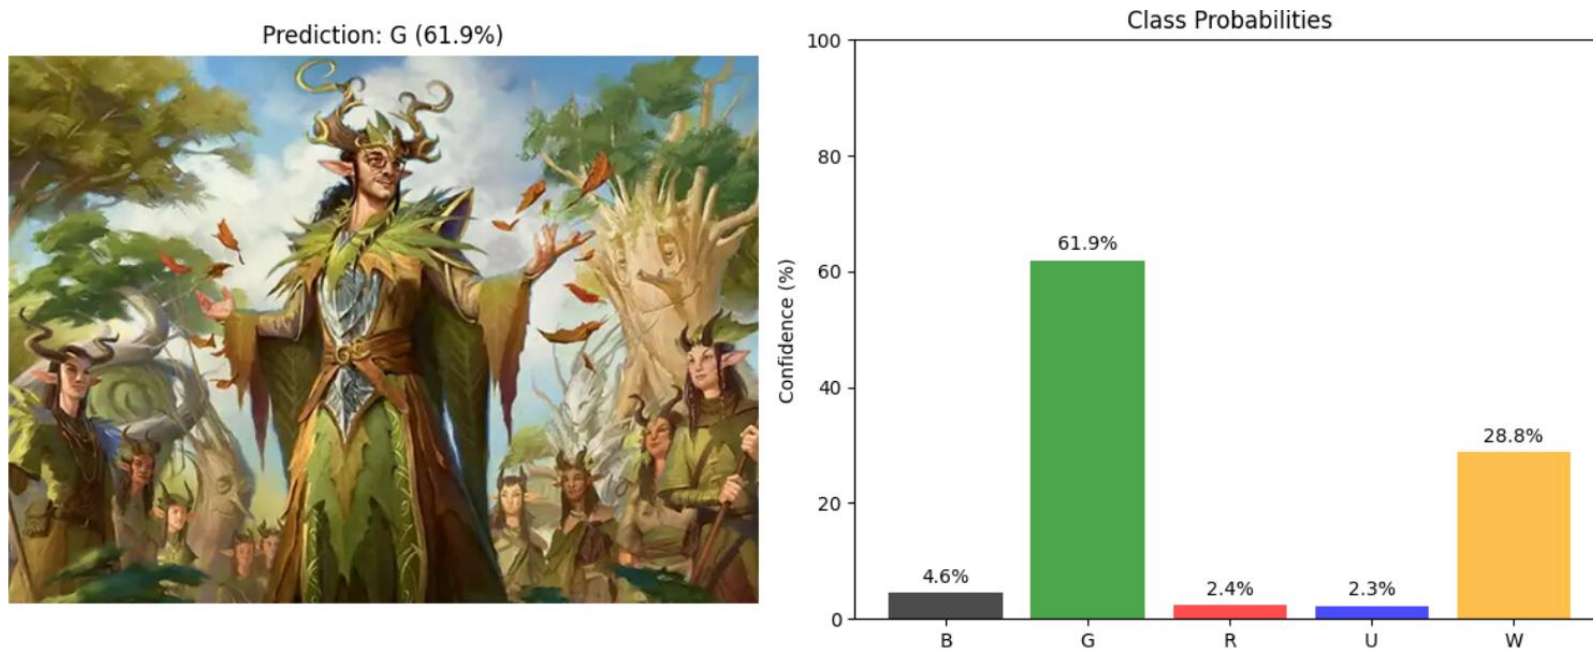
Quick Reminder About Classification

- Classification models don't just output a guess, they output a probability distribution across all classes.
- Example:




Quick Reminder About Classification

- Models don't just output a guess, they output a probability distribution across all classes.
- Example 2:




Correct Classifications

- We can look at the images our model classified correctly with the highest confidence (assigned the most probability to the correct class)
- What does this tell us about our model?



Correct Classifications

Confusion Matrix



Color: B



Color: G



Color: R




Color: U





Color: W




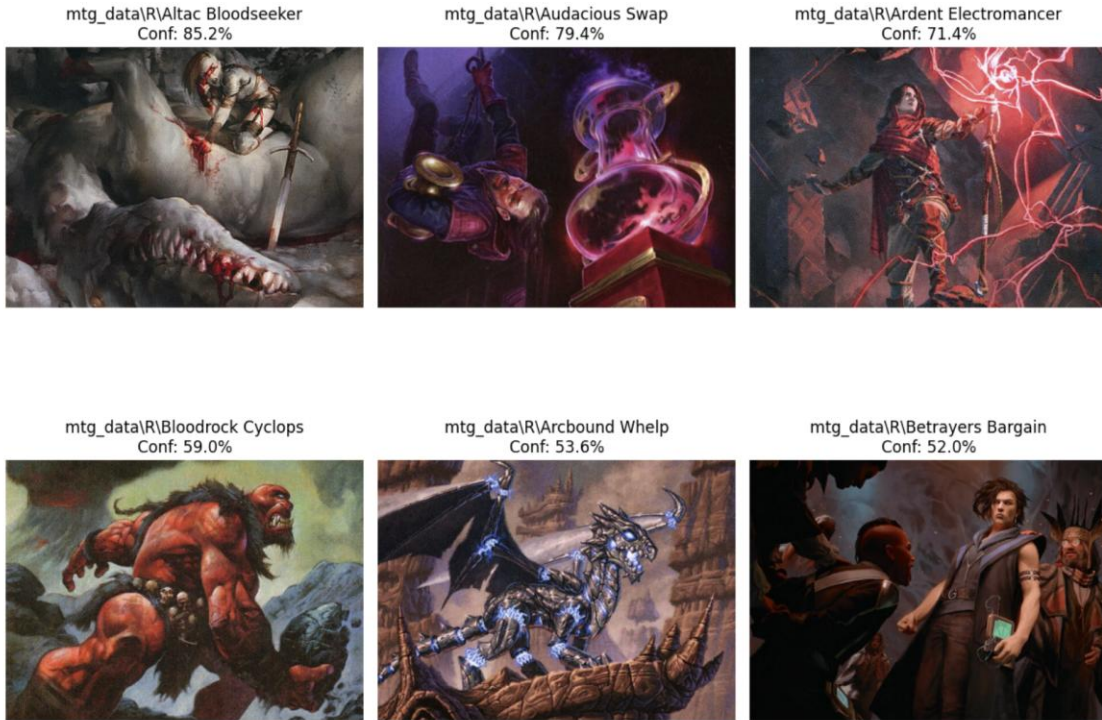
Confusion Matrix Revisited



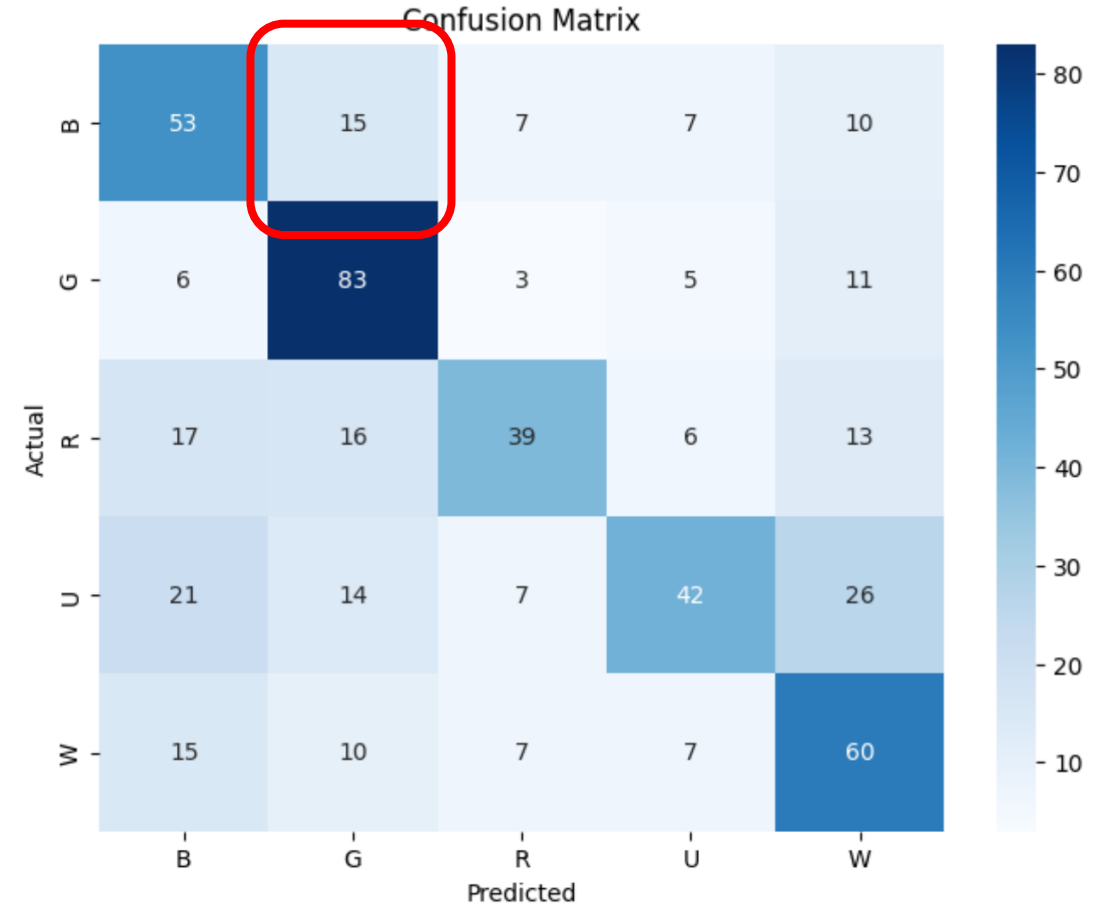

Confusion Matrix Revisited



Confusion Matrix Revisited



Confusion Matrix Revisited



Test the Model – Confusion matrix

- What does the confusion matrix tell us about how our model is doing?
- How did inspecting individual examples help us understand our confusion matrix/model?

