

From Zero to Neural Networks: A Structured Learning Guide for fast.ai Lesson 3

Integrating Jeremy Howard's fast.ai Lesson 3 with proven learning science

Dr. Neal Aggarwal | drnealaggarwal.info

Based on Jeremy Howard's fast.ai Practical Deep Learning for Coders 2022 ·
Lesson 3: Neural Net Foundations · Supplementary Swadia Tutorial
(youtu.be/npQ2IORdlvU)

About This Guide

This structured learning guide accompanies fast.ai Practical Deep Learning for Coders Lesson 3: Neural Net Foundations. It has been prepared by Dr. Neal Aggarwal (drnealaggarwal.info) — an educator with 40+ years of experience teaching information technology and AI — to help learners move from surface-level comprehension to deep, transferable understanding of how neural networks work.

The guide integrates Jeremy Howard's pedagogical approach, established learning science on memory retention and skill acquisition, and supplementary material from the Swadia tutorial video (<https://youtu.be/npQ2IORdlvU>). Read it alongside the NotebookLM resource <https://notebooklm.google.com/notebook/8e7184c3-a0c5-47a5-8992-85e90f8b5261> at

Lesson 3 at a Glance: The Core Argument

Jeremy Howard's Lesson 3 asks a deceptively simple question: 'What are the weights in a trained model, and how can numbers figure out something important about the world?' The entire lesson is a rigorous answer, built from first principles.

The core argument has five moves:

1. Any function can be parameterised — exposing free parameters turns 'find the right function' into 'find the right numbers'.
2. Loss converts dataset-level error into a single differentiable scalar.
3. Gradients are the local slope of the loss — PyTorch computes them with `.backward()`.
4. Gradient descent is the algorithm — subtract a scaled gradient, repeat.
5. ReLUs are universal approximators when stacked — enough simple piecewise-linear functions approximate any continuous mapping.

Swadia Video: 10 Key Learning Points

These points are drawn from the supplementary tutorial at <https://youtu.be/npQ2IORdlvU> and structured as an ordered learning sequence complementing Howard's top-down approach.

POINT 1 — Context First, Mechanics Second Before any equation, establish why it matters. Neural networks exist to solve the function-approximation problem. Anchor every

formula in this problem statement.

POINT 2 — The Biological Metaphor Is a Scaffold, Not a Specification The McCulloch-Pitts neuron (1943) gave us the vocabulary: weighted inputs, threshold, output. Treat this as intuition-building scaffolding, not a literal description.

POINT 3 — Parameters Are the Entire Knowledge of the Model After training, a model IS its weight matrices. There is no other store of information. This dispels the mysticism around 'AI knowing things'.

POINT 4 — Loss Is a Design Choice MSE, cross-entropy, Huber loss — these are engineering choices, not revealed truths. Practice swapping loss functions and observe the effects.

POINT 5 — The Learning Rate Is the Most Important Hyperparameter Too large: loss diverges. Too small: training takes forever. Learning rate scheduling follows from this single observation.

POINT 6 — One Hidden Layer Is Enough; Depth Adds Efficiency A single hidden layer can approximate any function. Depth buys efficiency: the same approximation using far fewer parameters. This is why 'deep' learning works.

POINT 7 — Overfitting Means the Model Learned the Training Data, Not the Task Train loss down, validation loss up — that is the diagnostic. Regularisation, dropout, and data augmentation are interventions at this diagnostic point.

POINT 8 — GPUs Accelerate Matrix Multiplication, Not Magic A forward pass is a sequence of matrix multiplications interleaved with elementwise nonlinearities. GPUs are designed precisely for this pattern.

POINT 9 — Feature Engineering Is Still Required Log-transforming skewed features, normalising inputs, and dummy-encoding categoricals are not optional. They directly affect the condition number of gradient updates.

POINT 10 — Transfer Learning Is the Practical Default Training from scratch is rarely necessary. Pretrained models encode general feature hierarchies that fine-tune to new tasks with minimal data and compute.

Learning Science: How to Make This Stick

SPACED REPETITION Return to concepts after increasing intervals (1 day, 3 days, 1 week, 2 weeks). Use Anki for vocabulary (gradient, loss, epoch, batch, learning rate) and code patterns.

INTERLEAVED PRACTICE Mix gradient descent problems with ReLU visualisation with spreadsheet work. Interleaving increases difficulty and dramatically improves retention.

RETRIEVAL PRACTICE Close the notebook. Write the gradient descent loop from memory. Retrieval is not testing — it IS learning. Every correct recall strengthens the memory trace.

ELABORATIVE INTERROGATION For every concept, ask 'why?' not just 'what?'. Why subtract the gradient? Why `requires_grad=True`? Why `grad.zero_()`? These build causal understanding.

WORKED EXAMPLES BEFORE PROBLEM SOLVING Study Howard's notebooks before attempting independent implementation. Cognitive load theory shows novices learn more from worked examples than from struggling alone.

DESIRABLE DIFFICULTY Howard's top-down approach creates desirable difficulty. The discomfort of encountering the whole before understanding the parts is the signal that learning is occurring.

DUAL CODING Draw the gradient descent bowl. Sketch the ReLU function. Plot the loss curve by hand. Visualising forces commitment to a specific representation.

TEACHING OTHERS The Feynman Technique: explain the concept as if to someone with no background. Every confusion in the explanation is a gap in your own understanding.

Step-by-Step Study Plan

WEEK 1 — FOUNDATION

Day 1: Watch Lesson 3 video with note-taking. Capture questions, not answers. <https://www.youtube.com/watch?v=hBBOjCiFcuo>

Day 2: Run the Kaggle notebook. Execute every cell. Change parameters. Observe effects. <https://www.kaggle.com/code/jhoward/how-does-a-neural-net-really-work>

Day 3: Watch the Swadia supplementary video. Capture his 10 key points in your own words. <https://youtu.be/npQ2IORdlvU>

Day 4: Attempt the Titanic spreadsheet from scratch. Do not copy — struggle productively.

Day 5: Re-watch from the 54-minute mark (matrix multiplication). Implement the two-layer network in a blank Jupyter notebook without reference.

Day 6: Explore the NotebookLM audio overview. Listen while following your notes. <https://notebooklm.google.com/notebook/8e7184c3-a0c5-47a5-8992-85e90f8b5261>

Day 7: Post a question or explanation on forums.fast.ai.

WEEK 2 — CONSOLIDATION

Day 1: Re-implement the gradient descent loop from memory. Day 2: Submit a Kaggle Titanic result. Day 3: Read Chapter 4 of the fastai book. Note what the video compressed. Day 4: Record a short video re-teaching the lesson for yourself. Day 5: Begin Lesson 4 (NLP), carrying the conceptual framework forward.

CHECKPOINT QUESTIONS

Before moving on, answer these without looking:

- Why does ReLU enable universal approximation when linear layers alone cannot? - What does `params.grad.zero_()` do, and what happens if you omit it? - Why is log-transforming Fare better than leaving it raw? - What is the difference between training loss and validation loss? - How do you know if your learning rate is too large or too small?

Code Patterns to Memorise

These are the minimum viable code patterns for Lesson 3. Write each from memory.

1. GRADIENT DESCENT LOOP

```
params = torch.tensor([a, b, c], requires_grad=True)
for _ in range(epochs):
    preds = model(params, x)
    loss = mse(preds, y)
    loss.backward()
    with torch.no_grad():
        params -= lr * params.grad
    params.grad.zero_() # <-- do not forget this
```

2. RELU ACTIVATION

```
def relu(x): return x.clamp(min=0)
```

3. LINEAR LAYER

```
W = torch.randn(n_in, n_out, requires_grad=True)
b = torch.zeros(n_out, requires_grad=True)
def linear(x): return x @ W + b
```

4. TWO-LAYER NEURAL NETWORK

```
def model(x): return linear2(relu(linear1(x)))
```

5. MSE LOSS

```
def mse(preds, targets): return ((preds - targets)**2).mean()
```

What AI Brings to Learners Who Act Now

The AI transition is not a future event — it is underway. Learners who invest time in foundations now gain access to capabilities that currently require expensive specialists:

PERSONAL PRODUCTIVITY TOOLS Fine-tuned language models for your domain. Image classifiers trained on your own data. NLP extractors that pull structured information from unstructured text.

PROFESSIONAL DIFFERENTIATION In every data-intensive field — medicine, law, finance, engineering, research — practitioners who build and evaluate AI tools command significant advantages.

RESEARCH ACCELERATION Custom models reduce dependence on off-the-shelf tools that may not fit your domain. FastAI puts state-of-the-art techniques within reach of any determined practitioner.

AGENCY IN AN AI WORLD Understanding how neural networks work gives you the conceptual tools to evaluate AI claims critically, audit model outputs, and make informed decisions about when to trust and when to question AI systems. This is the difference between using AI and being used by it.

About Dr. Neal Aggarwal

Dr. Neal Aggarwal (drnealaggarwal.info) has been teaching information technology and artificial intelligence for more than 40 years. His approach combines rigorous technical content with practical application and genuine pastoral care for each learner's progress.

Dr. Neal offers:

- One-to-one tutoring sessions (fast.ai and beyond)
- Group workshops for organisations introducing AI capability
- Curriculum design for academic and corporate AI education programmes
- Ongoing mentorship for practitioners transitioning into AI roles

Contact and booking: <https://drnealaggarwal.info>

The `fastai` book is available on Amazon: <https://www.amazon.com/Deep-Learning-Coders-fastai-PyTorch/dp/1492045527>

The entire book is also available free on GitHub for those who cannot afford to buy it: <https://github.com/fastai/fastbook>

