LESSON 1

LISTENING

Activity 1. You're going to watch an extract from an interview with a Googler about **machine learning**.

What do you already know about the topic? Have you ever worked on a project that used machine learning?

Activity 2. Watch the extract once. Link: https://www.youtube.com/watch?v=IcDV9ZjQ6qk (up to 1:23).

1. What is the main difference between **machine learning** and 'regular' programming?
2. What other details can you catch? Take notes and then share what you caught with your partner.

Notes:

Activity 3. Listen to this extract again. Your teacher is going to play each line several times. Write down what you hear.

0:19 I'm Laurence Moroney, _____ with Joshua Gordon.
0:23_____
_____ what machine learning is in a concrete manner.
0:28_____
_____ manual rules.
0:32 LAURENCE MORONEY: _____
0:34 JOSHUA GORDON: _____
0:35 _____
0:38 _____
0:39 _____
0:41 _____.
0:42 LAURENCE MORONEY: _____
0:43 JOSHUA GORDON: From examples.

Activity 4. In pairs/small groups, summarize the **features of connected speech** that made the extract from Task 3 difficult to transcribe:

Activity 5. Your teacher is going to play more sentences that contain some of the words difficult to catch.
First time you hear a sentence, simply try to catch that word. If you've caught the word, try to catch the words before and after it.

Activity 6.
Work in pairs. Cover the transcript below with a piece of paper. Listen to the rest of the interview line by line. Each time

(a) Listen to a line and try to catch what the speaker says (**concentrate on the meaning:** there's no need to remember the exact words at this stage).
(b) Share with your partner what you caught.
(c) Listen to the line once again and right after that uncover the line, read it and underline the parts that you failed to catch.
(d) Play the line in your head, preparing to hear it again.
(e) Listen to the line one more time, without reading it. Did you hear everything?

0:43 LAURENCE MORONEY: So pattern matching. It might be visual, or it might be other patterns that are hidden in data.
JOSHUA GORDON: Absolutely.
0:48 And so the input to machine-- so the beauty of machine learning, and the real secret sauce, is that an algorithm that learns patterns from data can solve thousands of different problems.
0:58 And the reason is if I write a Python program to recognize digits, my program is hard coded to work with digits.
LAURENCE MORONEY: Got it.
1:04 JOSHUA GORDON: But if I write an algorithm to learn patterns from data, I can use that for speech recognition, image recognition, medicine.
1:11 Basically, anything that you can start with examples, just tell apart A and B, my same algorithm that I wrote just once can tackle all these problems.
1:20 And that's a really special and actually fairly profound thing.

Activity 7. Sit back, close your eyes and listen to this part of the interview again.

Activity 8. Watch the next part of the interview (https://youtu.be/IcDV9ZjQ6qk?t=1m23s up to 4:07).

1) What is 'gradual descent'? (1:23 up to 3:07)
2) What are the applications of gradual descent? (3:07 up to 4:07)
3) What other details can you catch? Share what you caught with your partner.

Notes:

[Optional homework]

Ss shadow the first several sentences of Joshua Gordon's monologue.

LESSON 2

EXPLAINING COMPLEX IDEAS: structuring your speech

Warmer. Work individually. Brainstorm 8-10 **terms**, **tools** or **approaches** that are used in your profession that your groupmates probably don't know.

•	•	•
•	•	•
•	•	•

Work in pairs/small groups. Choose one term from your partner's list and ask them to explain it to you.

A 1. Look at the sentences underlined in Extracts 1, 2 and 3 below.
In which sentences does the speaker

(1) give a 'high level' explanation of a concept by reformulating it in simpler, non-professional terms? ___, ___, ___
(2) give examples? ___, ___, ___
(3) summarize the ideas that he has already mentioned? ___, ___, ___

Extract 1.

LAURENCE MORONEY : Now, it's something that a lot of people don't really understand what machine learning is in a concrete manner.

JOSHUA GORDON: (A) So machine learning is about learning from examples rather than writing manual rules. So the short way of saying that is regular programming is you write a lot of manual rules to solve a problem. In machine learning, you let the algorithm find those rules for you. And so the input to machine-- so the beauty of machine learning, and the real secret sauce, is that an algorithm that learns patterns from data can solve thousands of different problems. And the reason is if I write a Python program to recognize digits, my program is hard coded to work with digits. But if I write an algorithm to learn patterns from data, I can use that for speech recognition, image recognition, medicine. **(B)** Basically, anything that you can start with examples, just tell apart A and B, my same algorithm that I wrote just once can tackle all these problems. And that's a really special and actually fairly profound thing.

Extract 2.

JOSHUA GORDON: And instead of equations, we're going to use visual examples. So an equation could be like if you talk about gradient descent, **(C)** gradient descent basically means finding the minimum of a function. So if I just say that, like as a developer, I'm like, all right, what does that mean? **(D)** So you can think of any equation, like x cubed plus y squared plus whatever equals 7. There's some value of x and y. Anyway, you can find the bottom of that curve literally – **(E)** think of it as a bowl. You can drop a piece of fruit in a bowl and it will roll to the bottom. **(F)** And gradient descent just means finding where this function is 0.

Extract 3.

So in machine learning, let's say you're writing an algorithm. **(G)** Let's say it's to distinguish apples from oranges. You always want to know, how accurate is my algorithm? **(H) Like**, I can solve that problem in one line. I can just say, return math.random. So one line, math.random. My accuracy is crap. But really, we want to get-- another way of describing accuracy is you can think about it in terms of error. High accuracy means low error. And you can have an equation that describes your error. And the minimum of that equation is going to give you the highest accuracy. **(I)** So you can write your machine learning algorithm to try and minimize the equation that describes the error.

Task 2. Organize the linkers from Extracts 1, 2 and 3 in the table:

	Linking expressions:
Give a 'high level' explanation in simple terms	
Give an example(s)	
Summarize / reformulate	

Task 3. Choose **three terms** that you brainstormed during the warmer and write explanations using the structure and linkers from **Task 2.**

1.

2.

3.

Key

Task 1.

1. High level explanation: A, C
2. Examples: D, E (metaphor), H, G
3. Summary: B, F, I

Task 2.

	Linking expressions:
Give a 'high level' explanation in simple terms	So, [ ] **is about** [v-ing] [gradient descent] **basically means**…
Give an example(s)	**So you can think of any** [equation], **like** [x squared plus y], **whatever.** **Let's say** you're writing an algorithm. **Let's say** it's to distinguish apples from oranges. **Like**, I can solve that problem in one line. **[you can] think of it as** a [bowl]
Summarize / reformulate	**Basically**, this is any algorithm that starts from examples. **And** gradient descent **just means** [finding where…] **So**, [you can write your machine learning algorithm to try and minimize the equation that describes the error]

Homework.

1. Revise the expressions for explaining complex ideas using a Quizlet set:

https://quizlet.com/363791251/explaining-complex-ideas-in-simple-terms-flash-cards/

2. Rehearse explaining an idea/term that others in the group are unlikely to know using the framework and linkers that we studied in class

Success criteria:
I structured my speech to make it easy to follow [High level explanation / examples / reformulation or summary]
In each of the three parts of the monologue, I used some of the linkers we've studied, to signpost the structure
I used intonation and sentence stress to highlight important bits

LESSON 3

Students quietly rehearse their monologues for a few minutes and then mingle to explain the terms to those people who don't know them.

T circulates to provide on-the-spot and delayed feedback on accuracy.

Homework

Below is an extract from a Hangout in which Google employees help developers troubleshoot their problems.

(1) Watch the video once to get a basic idea of what the speakers are talking about https://youtu.be/E7V1xadkr-c?t=15m25s (up to 17:50)
(2) For questions (1)-(15), choose the best option to fill the gap from the table below.
(3) After you finish, watch the video again and check your answers. https://youtu.be/E7V1xadkr-c?t=15m25s

Android Developer Office Hours

>>Al: I get the impression from Ravi that one of the things that's providing a bit of confusion is the difference between densities of assets and screen sizes. It's best if you (1)_____ two different things.

(2)_____, the density of assets (3)_____ how much detail you want. And the screen size is your layout that you want to incorporate. (4)_____ that way, when you come across (5)_____ an S3, it's using the high-density assets but it's still using the right layout for its size screen.

>>Ravi: Right.

>>Sparky: Yeah, exactly. (6)_____, to sort of riff on what Al is saying, (7)_____, you have, say, a visual thing (8)_____ a button background or an icon. And you say, "I want this thing to be about two centimeters across." And then, you design a bitmap that measures about two centimeters across in each of (9)_____, three or four bitmap or pixel densities.

(10)_____, you say this thing is at 300 dots per inch or 200 dots per inch or 160 dots per inch. This is how many pixels across the image needs to be in order to measure (11)_____, two centimeters in real distances. (12)_____, you just have different versions of every visual asset and let the Android device know its own screen density and say, OK, I am an XDPI, XHDPI device.

I will pick this one. Or, I am an NDPI device. I'll pick that one. And you can then trust that when that thing has gone on screen, that is approximately two centimeters across, subject to minor variations because of the different size buckets. And then, so you size your assets for the density and you size your layouts for the actual size.

(13)_____ you say, I want, this is the layout I want to use and this screen actually physically measures seven inches across, or ten inches across, (14)_____. And then rely

on the device to pick and choose the right combination of layout per its screen size and the individual asset elements based on its pixel density. (15)_____.

1	A say	B think of those as	C basically
2	A Like	B Say	C So
3	A is about	B basically	C say
4	A So	B Think of it as	C say
5	A basically	B something like	C whatever
6	A Say	B So	C whatever
7	A whatever	B and so	C basically
8	A like	B think of it as	C and so
9	A whatever	B say	C and so
10	A Think of it as	B Something like	C So
11	A basically	B is about	C say
12	A Something like	B whatever	C And so
13	A So	B Whatever	C Did that make sense
14	A about	B something like	C or whatever
15	A Basically.	B Did that make sense?	C Say.

Source: Android Developers Youtube Channel.