GMU

GULF MEDICAL UNIVERSITY

College of Pharmacy

OSCEva
OSCE Assessor

The complete user manual — how to set up, transcribe, review, score, moderate and report an OSCE station, screen by screen.

Assessment Coordinator | Examiner | Assessment Committee

API transcription edition
Examiner-verified OSCE scoring

Application build 2026-10-26 · V2 Moderation
Third edition

What is in this manual

The manual follows the five steps of the application, in the order you use them. The first chapter is worth ten minutes whoever you are.

EVERYONE

1 How the application works 2

The five steps · what happens to a recording · data · the screen · phones and tablets · words

IF YOU SET UP THE STATION

2 Setting up 7

Choosing a speech API · Gemini · OpenAI · Vertex AI · privacy and security measures

3 The station rubric 12

Adding · how the rubric is read · checking · levels · verifying · changing a rubric

4 Recordings and the roster 18

The roster · adding recordings · analysing · naming the students · more actions

IF YOU EXAMINE

5 Review and score 22

The review screen · identity · transcript · confidence and filters · items · feedback · sign-off

WHEN THE MARKING IS DONE

6 Reports 29

Report details · final mark · student report · cohort · exports · moderation · Cronbach's alpha

7 Reference 38

When something goes wrong · labels · privacy · good practice

CHAPTER ONE

How the application works

One idea explains almost everything else: the application **proposes** and the examiner **decides**. It listens to a recording, writes down what was said, finds each student and drafts a score for every rubric item — and then hands every one of those drafts to a person who checks it against the video and signs it.

Read this chapter **whatever your role**

The five steps

Every station moves through the same five steps at the top of the screen, left to right. A step can be revisited at any time, but each one needs the one before it.

STEP ONE

Setup

Paste the API key that transcribes the recordings.

STEP TWO

Rubric

Upload the station's marking rubric, check it against the source and verify it.

STEP THREE

Recordings

Load the class roster and the OSCE videos, then analyse them.

STEP FOUR

Review & score

Check every proposed score against the video, edit, and sign off.

STEP FIVE

Reports

Enter the report details and exam mark, print reports and export results.

THE RULE THAT SURPRISES PEOPLE

The AI does **not** mark the students. The speech API is used to turn the sound into a transcript and to say who spoke each line. Finding the students and matching every rubric item against what the **student** said is done on your own computer, by rules you can read on screen, and nothing becomes a mark until an examiner has reviewed it and signed it off.

One recording, many students

A single video may hold a whole session — student after student at the same station. The application finds where each student starts and stops, cuts the recording into one section per student (the file itself is never changed), and gives each student their own record, review and report.

What happens to each recording

When you press **Analyse**, the application works through these steps on its own. The live log on the Recordings step shows each one as it happens.

1. **Listen.**
   The sound is taken from the video and cleaned on this computer. Long silences are skipped.
2. **Transcribe.**
   The sound — never the picture — is sent in pieces of up to two minutes to Google Gemini, Google Vertex AI or OpenAI, which return the words. Each word then gets its exact time from the sound itself.
3. **Check the words.**
   Two more speech models hear the same pieces. Where two of three agree on a different word it is corrected; where all three differ it is marked for you.
4. **Mark who spoke.**
   Each line of the transcript is labelled Student, Patient or Examiner.
5. **Find the students.**
   The transcript is split into one dialogue per student, and each student is recognised by the name — then the ID — they say, matched against the class roster.
6. **Take pictures.**
   Pictures of every student are taken from their own section of the video, so that you can confirm who they are.
7. **Draft the scores.**
   Every rubric item is matched against what the student said — never the patient — with the time of each matching line as evidence and a confidence percentage.

Where your data goes

WHAT	WHERE IT GOES	KEPT
The sound of each recording	The speech API you chose (Google Gemini, Google Vertex AI or OpenAI), in pieces, with no file name, date or other metadata	Not stored by the app; any upload is deleted when the analysis ends
The transcript text	The same API, once, to mark who is speaking	Not stored by the app
The video picture	Nowhere — it stays on this computer	—
The class roster (names, IDs)	Nowhere — matching is done on this computer	—
Rubric, scores, comments, reports	Nowhere — this browser only	Until the tab is closed

The screen

Every screen shares the same frame: the OSCEva icon and name on the left, the five steps in the middle and two buttons on the right. A green dot on a step means it is ready; a number shows how many items wait there.

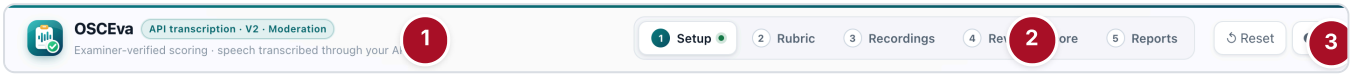


Figure 1.1 — The header with the new OSCEva icon. The step you are on is raised and numbered in colour.

1. **The application**
   The OSCEva icon, its name and edition.

2. **The five steps**
   Setup, Rubric, Recordings, Review & score and Reports.

3. **Reset and theme**
   ↺ Reset starts a new session (settings and key are kept); ◐ switches light and dark.

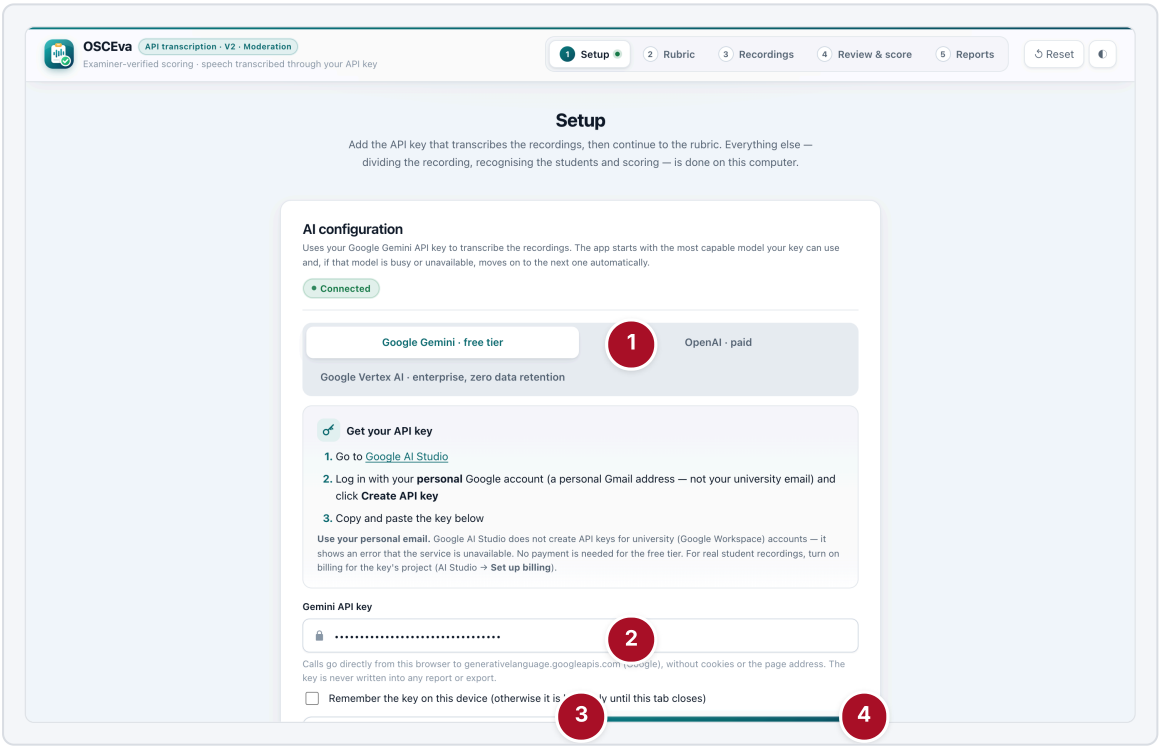


Figure 1.2 — The first screen: Setup. The key you paste here is the only thing the application needs before you start.

1. **Choose the speech API**
   Google Gemini (free tier), OpenAI (paid) or Google Vertex AI (enterprise).

2. **Paste the key**
   Hidden as you type and never printed on any report.

3. **Test connection**
   Checks the key and shows which models it can use.

4. **Connect & continue**
   Checks the key once more and opens the Rubric step.

On phones and tablets

The application works in the browser of a phone or tablet as well as on a computer. Every screen rearranges itself for the smaller width: the five steps stay visible as one row of buttons, recordings and the cohort summary become cards, and in Review the students become a strip you swipe sideways. Windows open from the bottom of the screen.

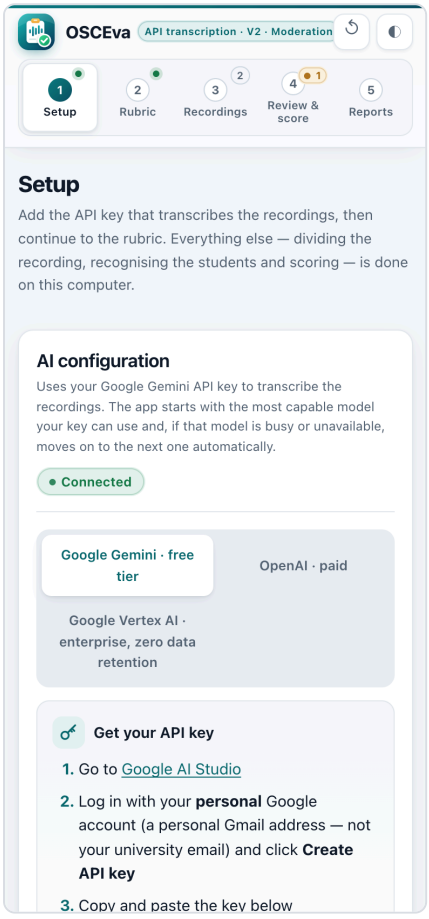


Figure 1.3 — Setup on a phone.

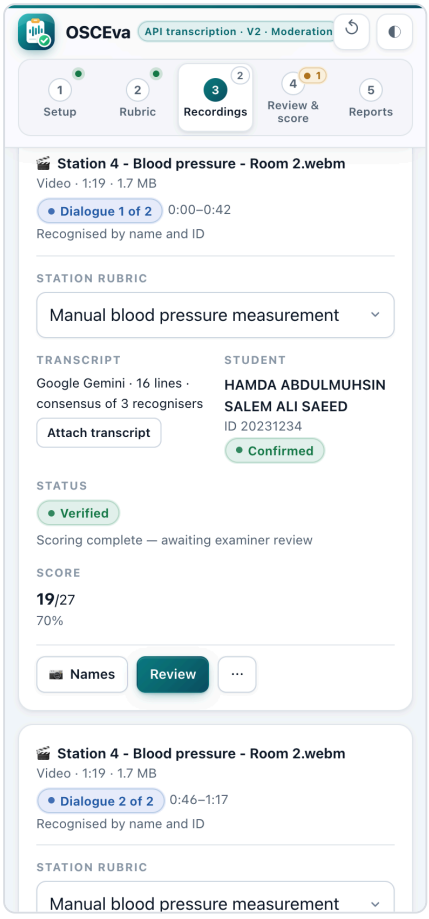


Figure 1.4 — A recording as a card on a phone.

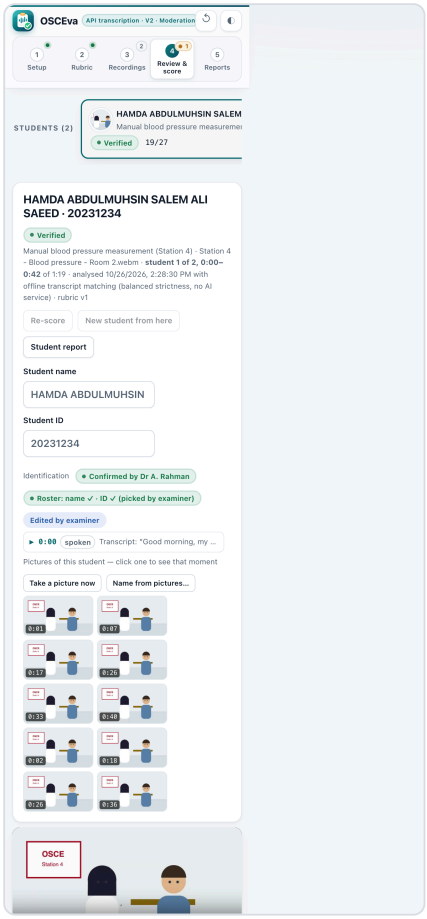


Figure 1.5 — Review on a phone: swipe between students.

ADD IT TO THE HOME SCREEN

On an iPhone or iPad, open the application in Safari, tap **Share → Add to Home Screen**; on Android, open it in Chrome and tap ⋮ → **Add to Home screen**. It then opens from the OSCEva icon like an app. Long recordings are best analysed on a computer; phones and tablets are ideal for reviewing and signing off.

Who does what

The application has no accounts or sign-in: whoever opens it can do everything. In practice the work divides between two people, and the chapters follow that split.

CHAPTERS 2 – 4

Assessment coordinator

Sets up the key, uploads and verifies the station rubric, loads the class roster and the recordings, and starts the analysis.

CHAPTER 5

Examiner

Confirms who each student is, checks every proposed score against the video, adds comments and signs off.

CHAPTER 6

Coordinator or examiner

Fills in the report details and exam mark, prints each student's report and exports the results.

CHAPTER 7

Everyone

What the labels mean, what to do when something goes wrong, and good practice.

Words used in this manual

WORD	WHAT IT MEANS HERE
Station	One OSCE task with its own marking rubric, for example "Manual blood pressure measurement".
Rubric	The marking checklist for a station: numbered items, each with its levels and points.
Level	One possible mark for an item: its points and a descriptor ("Done well — name, role and two identifiers").
Recording	One video or audio file. It may hold one student or a whole session.
Section	The part of a recording that belongs to one student, from their first word to their last.
Confidence	How well the student's words matched an item, as a percentage: Low, Medium or High.
Exam mark	The marks the station carries in the exam. The final mark is the score % of it.
Moderation	Comparing two (or three) graders' results for the same students and agreeing one result.
Cronbach's alpha	A number from 0 to 1 that says how consistently the items of a station measure the same performance.
Sign-off	The examiner's confirmation that the scores are their own judgement. Signed scores are locked.

CHAPTER TWO

Setting up

Setup is done once per computer. It asks for one thing — an API key for the speech service that writes the transcripts. Everything else is already set to the recommended values.

For the **assessment coordinator** — before the first analysis

Choosing a speech API

Three services can transcribe the recordings. All three are used directly from your browser with your own key; choose one at the top of the **AI configuration** card.

	GOOGLE GEMINI · FREE TIER	OPENAI · PAID	GOOGLE VERTEX AI
Cost	Free tier available; paid tier with billing	Pay as you go, about one US cent per minute	Your institution's Google Cloud billing
Word times	Phrase times from the API; each word timed from the sound	Every word timed by Whisper, refined from the sound	As Gemini
Three recognisers	Three Gemini models	Whisper, GPT-4o Transcribe and GPT-4o mini Transcribe	Three Gemini models your project can use
Account	Personal Google account	Personal email	Institution's Google Cloud project
Student privacy	Free tier: Google may use what you send — use a paid project for real students	Not used for training by default; kept up to 30 days	Google Cloud terms; **zero data retention** when the project is set up for it

USE YOUR PERSONAL EMAIL

Create the key with a **personal** account — a personal Gmail address for Google, a personal email for OpenAI — **not your university email**. Google AI Studio does not create API keys for university (Google Workspace) accounts: it shows an error that the service is unavailable in your region or for your account.

For real student recordings, Google Vertex AI is the option Google itself recommends for guaranteed zero data retention and enterprise data agreements (page 10).

Getting a Google Gemini key

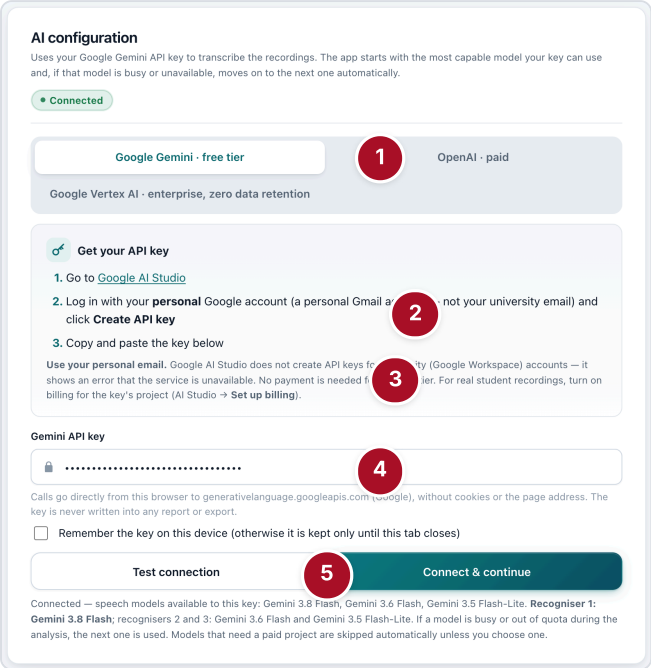


Figure 2.1 — The AI configuration card with Google Gemini chosen. Step 2 and the note beneath now ask for a personal Google account.

1. **Choose Google Gemini**
   At the top of the card.
2. **Sign in with a personal Google account**
   Open aistudio.google.com/apikey with your personal Gmail address — not your university email — and click **Create API key**.
3. **Why personal**
   University accounts are refused by Google AI Studio.
4. **Paste the key**
   It begins `AIza`.
5. **Test connection**
   The chip turns green: **Connected**, and the message lists the three recognisers.

What the test tells you

THE CHIP SAYS	WHAT IT MEANS
No key	Nothing has been pasted yet.
Key added · not tested	A key is there; press Test connection to check it.
Connected	The key works and at least one speech model is available.
Needs attention	The key works but a model is unavailable, or the network failed — read the message.
Key not working	The key was refused, or the account has no credit. Copy it again or check billing.

If a model is busy or its daily quota is used up during an analysis, the application moves on to the next model automatically and records which one it used.

Getting an OpenAI key

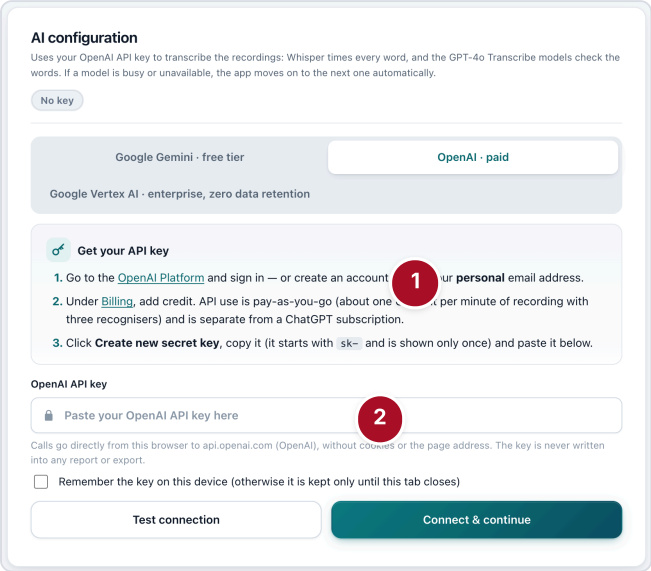


Figure 2.2 — The same card with OpenAI chosen.

1. **Sign in with a personal email**
   Open platform.openai.com/api-keys and sign in — or create an account — with your personal email address.
2. **Paste the key**
   Click **Create new secret key**, copy it (it starts sk- and is shown only once) and paste it. Add a little credit under **Billing** first.

> **ONE KEY PER SERVICE**
>
> Each service keeps its own key, so you can switch between Gemini, OpenAI and Vertex AI without pasting again. Switching is blocked while an analysis is running — press **Stop** first.

Options that work in the background

Setup now shows only the AI configuration. The transcription and analysis options are no longer on screen, but they still apply, with the recommended values:

OPTION	VALUE USED	WHAT IT DOES
Transcripts	Through the API	Every recording is transcribed; an attached transcript file is still used first.
Three speech recognisers	On	Every section is heard by three models and the words are put to a vote.
Recognisers	Automatic	The most capable models your key can use; the next one takes over if one is busy.
Spoken language	English	The language the models listen for.
Matching strictness	Balanced	How closely the student's words must match an item to earn credit.
One student per section	On	Faces are checked in every section, and a section is cut where a different face appears.

Connecting Google Vertex AI

Vertex AI runs the same Gemini models inside your institution's Google Cloud project, under the Google Cloud terms and your institution's data processing agreement. Google does not use the data to train its models, and the project can be set up for **zero data retention**.

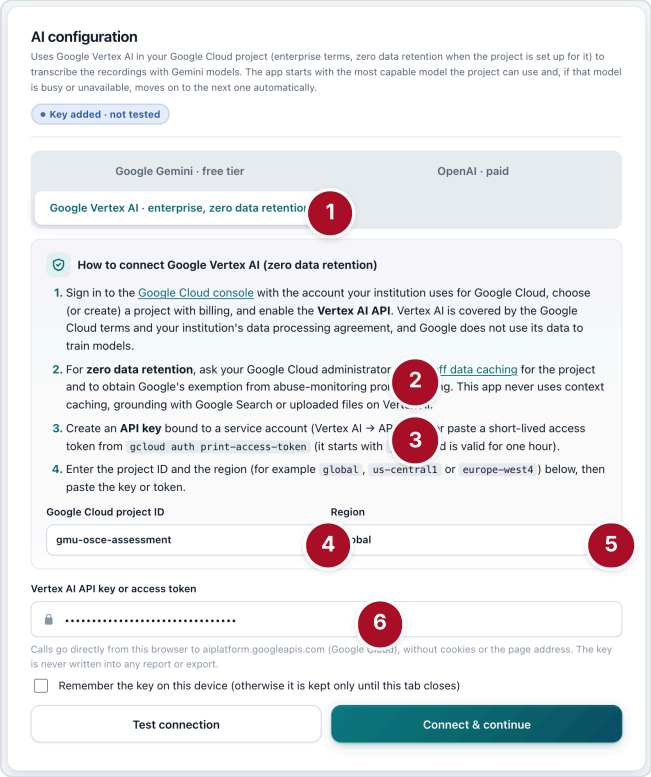


Figure 2.3 — The AI configuration card with Google Vertex AI chosen.

1. **Choose Google Vertex AI**
   The third option at the top of the card.
2. **Zero data retention**
   Your Google Cloud administrator turns off data caching for the project and obtains Google's exemption from abuse-monitoring logging.
3. **Key or token**
   An API key bound to a service account, or an access token (`ya29.…`, valid for one hour).
4. **Project ID**
   The Google Cloud project that is billed.
5. **Region**
   `global`, or a region such as `europe-west4`.
6. **Paste and test**
   Press **Test connection**: the models your project can use are listed.

WHO SETS IT UP

Vertex AI is normally set up once by the university's Google Cloud or IT team, who then give the examiners a key. An access token expires after an hour — paste a new one when the app says so.

Privacy and security measures

Whichever speech service you choose, the application protects the recordings in the same ways:

MEASURE	WHAT THE APPLICATION DOES
Delete as soon as the analysis ends	The sound is sent inside each request, not uploaded as a file. Should any file be uploaded to Gemini, it is deleted the moment the analysis of that recording finishes — whether it succeeded, failed or was stopped.
No context caching	The application never asks a service to cache a request. OpenAI requests are sent with storage turned off.
No metadata	The sound is re-encoded on this computer into plain audio with no embedded details, under a generic name. File names, file dates, the page address and browser cookies are never sent, and the original files' dates are not kept in backups or exports.
One service only	The application talks only to the service chosen in Setup. Every other AI or Google service is blocked inside the app.
Sound only	The video picture, the class roster, the rubric and the scores never leave this computer.
Zero data retention	Available with Google Vertex AI when the Google Cloud project is set up for it (page 10).

What the key is used for

The key is used only to send pieces of sound to be transcribed and the finished transcript once to label who is speaking. It is kept in this browser tab — or, if you tick **Remember the key on this device**, in this browser on this computer — and is never written into any report, CSV or backup.

> **YOUR INSTITUTION'S RULES STILL APPLY**
>
> Confirm that the service you choose meets your institution's data-protection policy and the students' consent before you assess real candidates.

CHAPTER THREE

The station rubric

The rubric is what every recording is scored against, so it must be exactly right. The application reads your rubric file into an editable checklist, and nothing can be scored until a person has compared it with the source and verified it.

For the **assessment coordinator**

Adding a rubric

Open **Rubric**. Drop the rubric file on the box, or click it to browse. Several files — a checklist and a global rating scale — can be combined into one station rubric.

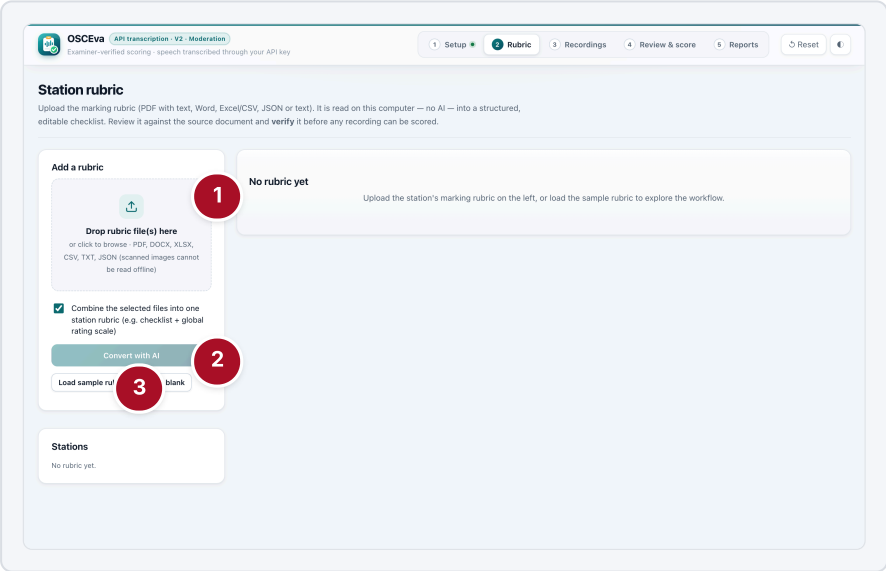


Figure 3.1 — The Rubric step before any rubric is added. **Load sample rubric** shows a complete example.

1. **Drop the rubric file**
   PDF with text, Word, Excel, CSV, JSON or plain text.

2. **Convert**
   Reads the file on this computer into items, levels and points.

3. **Sample or blank**
   Load the sample to practise, or start an empty rubric.

A PDF must contain text — a scanned picture of a page cannot be read. A rubric saved with **Export JSON** is read back exactly.

How the rubric is read

The reader now follows the rubric document exactly. It reads the marking table — wherever it is in the file and however many pages it runs over — and nothing else.

THE READER	IN PRACTICE
Ignores the case	The case scenario, patient details, instructions to the candidate and to the simulated patient are skipped. Only the marking rubric is read.
Reads every item	Every numbered item, from every page and every table of the file, in its own order — including tables whose header repeats on each page, merged cells and items without a number.
Reads levels exactly	Each item keeps exactly the levels, points and descriptors written in the file. No level, point or descriptor is added that is not in the rubric.
Keeps the descriptor whole	Lines and bullet points inside one table cell stay together in one descriptor.
Names the rubric	The rubric takes the name of the station activity in the file, for example "Manual blood pressure measurement".
Shows what it was unsure of	Anything it could not place is flagged for you to check in the editor.

> **ALWAYS COMPARE WITH THE SOURCE**
>
> However accurate the reader, the rubric is what every student is scored against: compare it item by item with the original document before you verify it.

Checking the rubric

The editor shows the station details at the top and every item below. Change anything that does not match your source document.

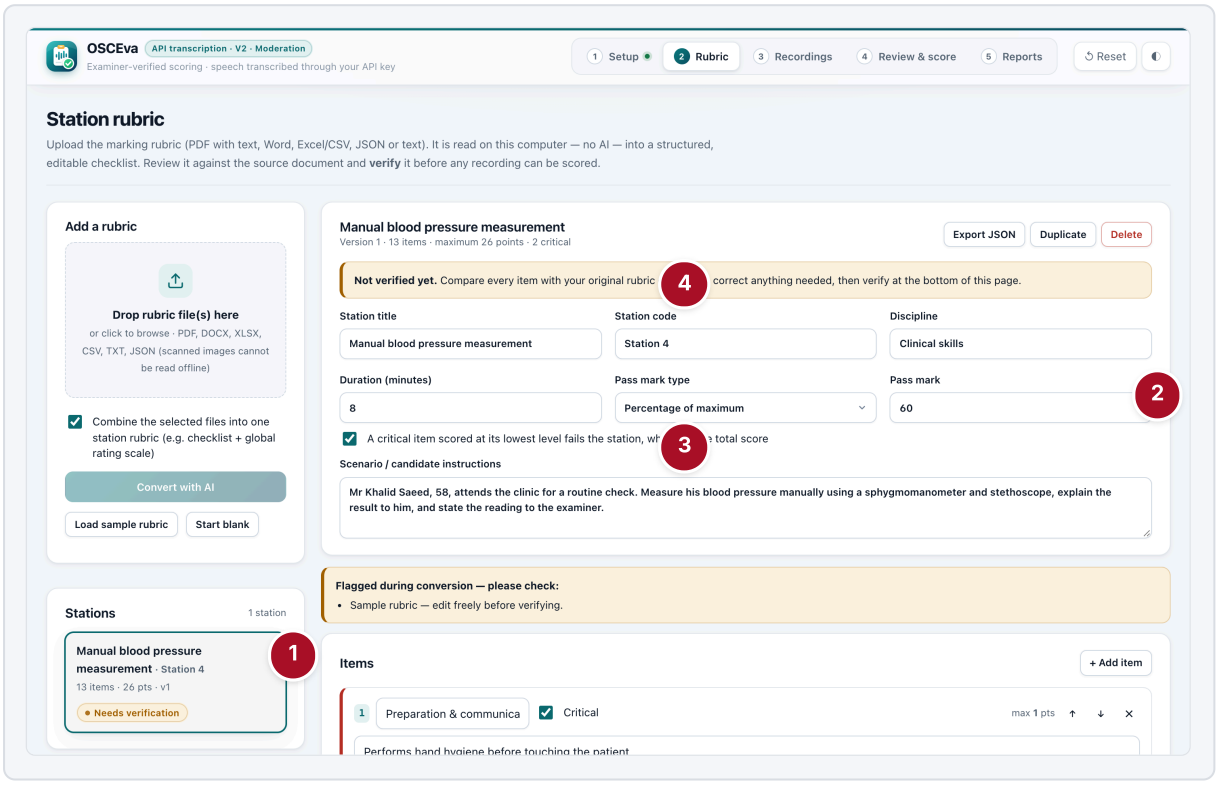


Figure 3.2 — The rubric editor: station details, pass mark and the critical-item rule at the top; the items follow.

1. **Stations**
   Every rubric in this session.

2. **Pass mark**
   A percentage of the maximum, a number of points, or none.

3. **Critical-item rule**
   A critical item scored at its lowest level fails the station.

4. **Flagged during conversion**
   Lines the converter was unsure about.

Levels: score and descriptor

Each item has its levels — the marks it can be given. A level is now just two things: its **points** and its **descriptor**. Where the source rubric gave a level a name ("Done well"), the name now opens the descriptor ("Done well — name, role and two patient identifiers"). To add a level, press **+ Add level** under the levels and type its points and descriptor. The item's maximum and the rubric's total update at once.

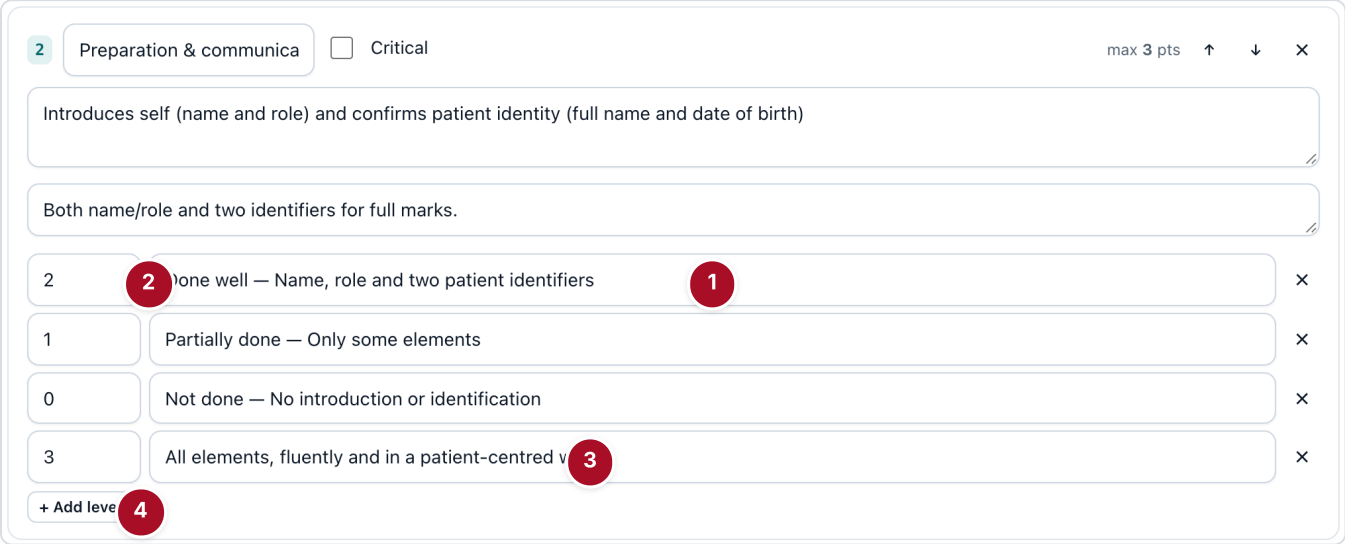


Figure 3.3 — An item with a new level worth 3 points. The item's maximum rose from 2 to 3.

1. **Descriptor**
   What the student did for this level, starting with the level's name.
2. **Points**
   Whole or half points.
3. **The new level**
   Points and descriptor of the added level.
4. **+ Add level**
   Adds another level to this item. ✕ removes one; an item keeps at least two.

> **LEVELS ARE SORTED BY POINTS**
>
> When the rubric is verified, each item's levels are put in order from the most points to the fewest, wherever you added them.

Verifying the rubric

At the foot of the editor is the verification box. Tick the three confirmations only after you have compared every item with the original document.

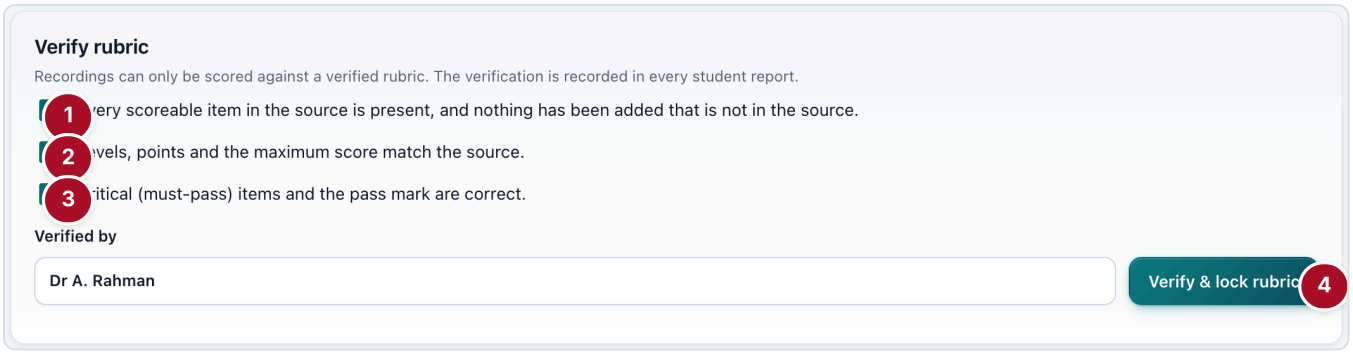


Figure 3.4 — Verify rubric. All three boxes and the verifier's name are required.

1 **Every item is there**
Nothing missing, nothing added.

2 **Levels and points match**
And the maximum score.

3 **Critical items and pass mark are right**

4 **Verify & lock rubric**
Records your name and the date on every report that uses it.

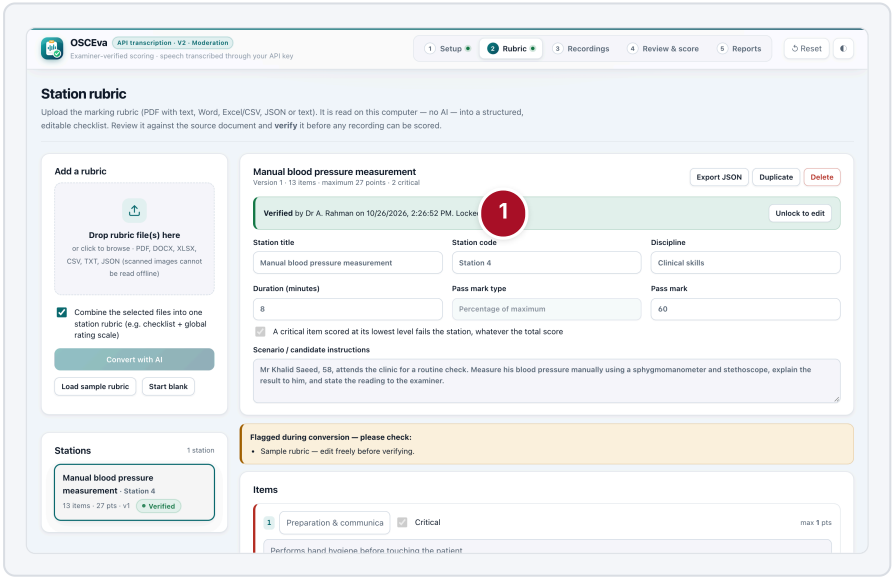


Figure 3.5 — A verified rubric. Its items are locked; recordings can now be scored against it.

Changing a rubric after scoring

A verified rubric is locked. To change it after students have been scored, press **Unlock to edit** on its green banner. The rubric becomes the next version and must be verified again.

1. **Unlock to edit**
   The rubric becomes version 2 (3, 4…) and opens for editing.
2. **Make your changes**
   Correct an item, add a level, change points or the pass mark.
3. **Verify again**
   Tick the three boxes and press **Verify & lock rubric**.
4. **Everyone is re-scored**
   Every student scored with the earlier version is re-scored with the new one at once. A message says how many.

RE-SCORING KEEPS	RE-SCORING REPLACES
The transcript, with who said each line	The proposed score of every item
Who the student is (their identity)	Examiner edits and comments on items
The recording and its sections	A sign-off — the student goes back to Awaiting review

> **NOTHING IS SENT AGAIN**
>
> Re-scoring uses the transcript already made, so it takes seconds and nothing is sent to the speech service. The previous result is kept in each student's history. A recording that is no longer loaded in the browser cannot be re-scored: re-attach it, then choose **Re-score** in its ⋯ menu.

CHAPTER FOUR

Recordings and the roster

Load the class roster, add the station videos, and press Analyse: each recording is transcribed, divided into one section per student, and scored — ready for an examiner to review.

For the **assessment coordinator**

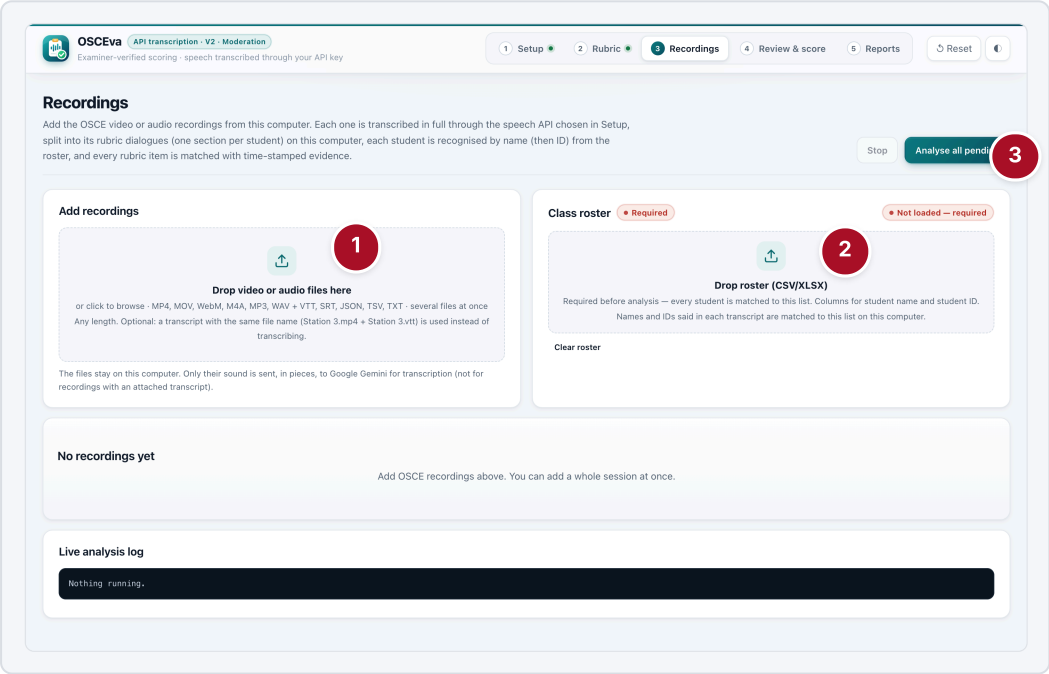


Figure 4.1 — The Recordings step: recordings on the left, the class roster on the right.

1. **Add recordings**
   Drop video or audio files — several at once, any length.

2. **Class roster**
   Required before any analysis.

3. **Analyse all pending**
   Starts every recording not yet analysed. **Stop** pauses the queue.

The class roster

COLUMN	EXAMPLE	NOTES
Student name	HAMDA ABDULMUHSIN SALEM ALI SAEED	Full name as on the university record
Student ID	20231234	The ID students say at the start of the station

Adding recordings

Drop the files on **Add recordings**. Each file becomes a row: its name, the station rubric it will be scored against, and its status.

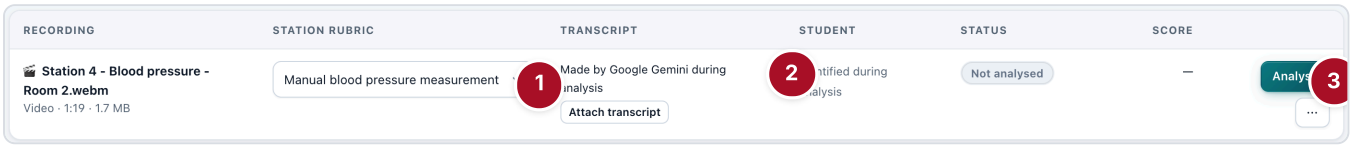


Figure 4.2 — A recording ready to analyse.

1. **Station rubric**
   Set automatically when only one verified rubric exists.
2. **Transcript**
   Made during analysis, or attach your own.
3. **Analyse**
   Starts this recording now; ⋯ opens more actions.

HOW A STUDENT IS RECOGNISED

Each section is read sentence by sentence. The student is recognised by the **name** they say, and by their **ID** only when no name is recognised. Names are transcribed exactly as spoken, and a name said by the patient is never taken for the student's.

Starting the analysis

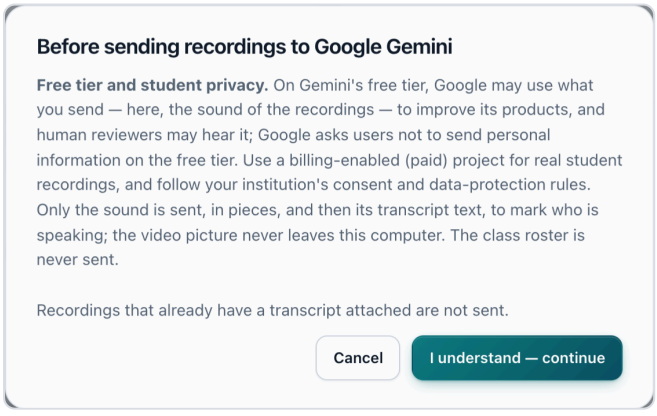


Figure 4.3 — The one-time confirmation before the first recording is sent.

Press **Analyse all pending**. The first time in each session you confirm that the recordings may be sent to the speech service. While a recording is analysed, its row shows the step it is on.

FIRST TIME ON A COMPUTER

The face-recognition models (about 17 MB) download once and are then kept by the browser.

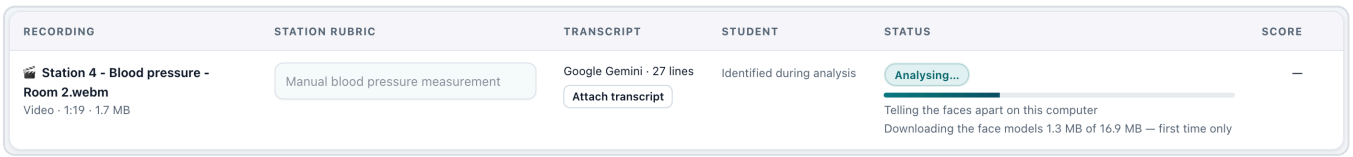


Figure 4.4 — A recording being analysed.

Naming the students

When the analysis ends, **Who is each student?** opens by itself. Each student has a card with pictures from their own section, and more pictures are taken automatically while the window is open.

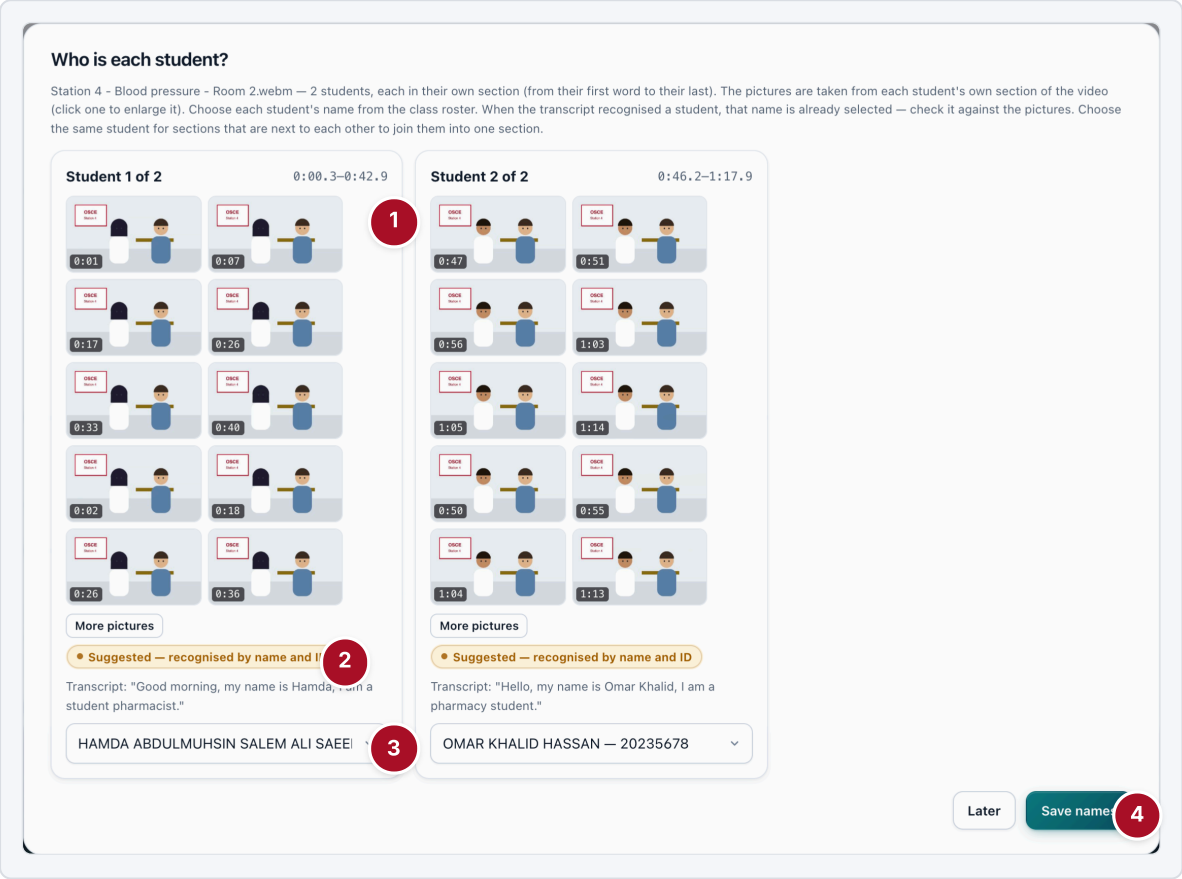


Figure 4.5 — Two students found in one recording, each recognised by name and ID.

1. **Pictures**
   From the student's own section. Click one to enlarge it.
2. **How they were recognised**
   By name and ID, by name, or not recognised.
3. **The roster name**
   Check it against the pictures and change it if needed.
4. **Save names**
   Confirms every student at once.

THE SAME STUDENT TWICE

If one student was split into two sections, choose the same name for both neighbouring cards: they are joined and scored again.

After the analysis

Each student now has a row of their own with the proposed score.

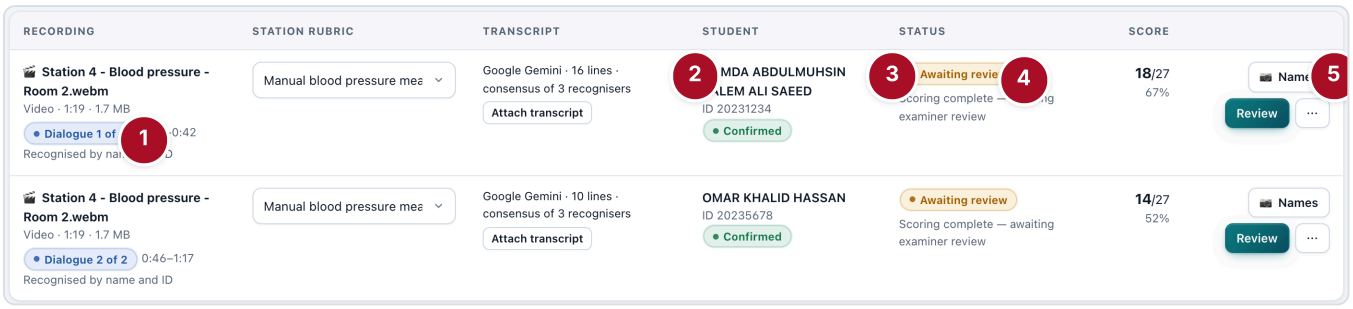


Figure 4.6 — Two students from one recording, each with their own dialogue, transcript, identity and score.

1. **Dialogue 1 of 2**
   Which section of the file this student is.
2. **Transcript**
   Which service made it.
3. **Student**
   Name, ID and confirmation.
4. **Status**
   Awaiting review until signed off.
5. **Names · Review · ⋯**
   Rename, review, or more actions.

Below the table, the **Live analysis log** lists every step of the analysis in plain words, with the time — the place to look when a result surprises you.

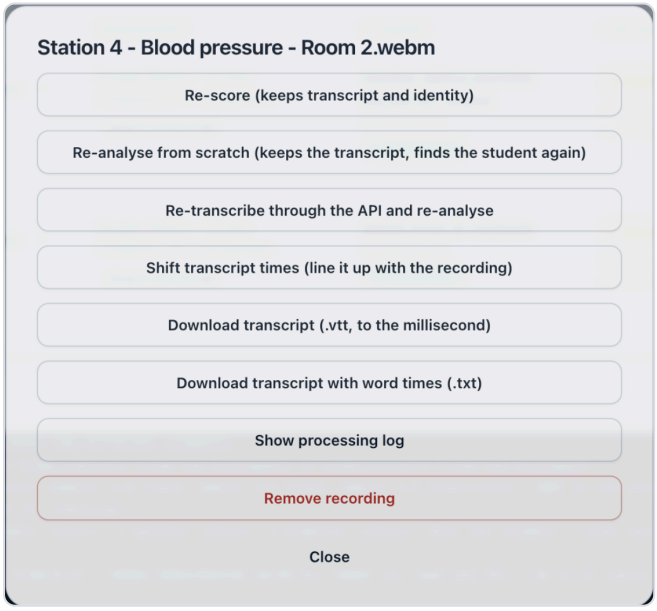


Figure 4.7 — The ⋯ menu of a recording.

ACTION	USE IT WHEN
Re-score	You changed the strictness. Keeps transcript and identity.
Re-analyse from scratch	Students were divided wrongly.
Re-transcribe	The transcript is poor.
Download transcript	As WebVTT, or text with word times.

CHAPTER FIVE

Review and score

Everything before this chapter produces drafts. Here the examiner turns them into marks: confirming who the student is, checking each proposed score against the video, adding comments, and signing off.

For **examiners**

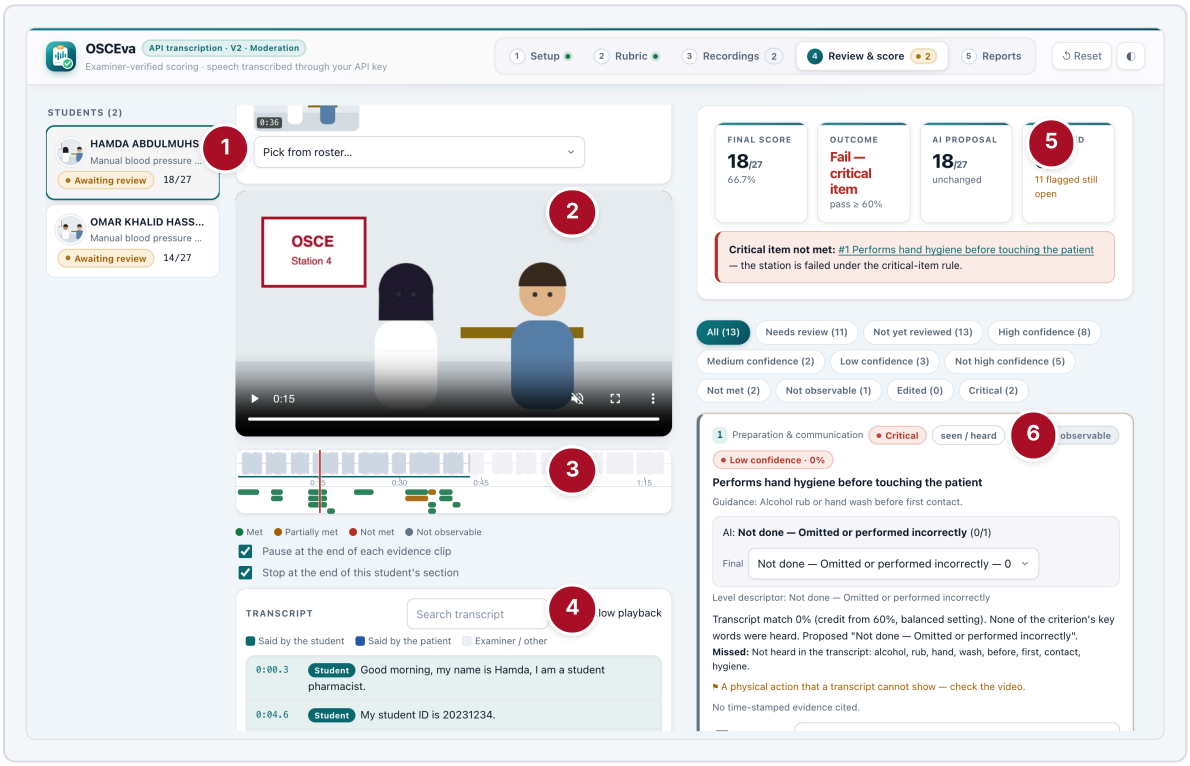


Figure 5.1 — The review screen, with the video paused at 0:15.

1. **Students**
   Everyone awaiting review.
2. **The video**
   This student's section only.
3. **Timeline**
   The sound and the evidence bars.
4. **Transcript**
   Who said what, word by word.
5. **Scores**
   Final score, outcome, AI proposal, reviewed.
6. **Rubric items**
   One card per item.

Who is this student?

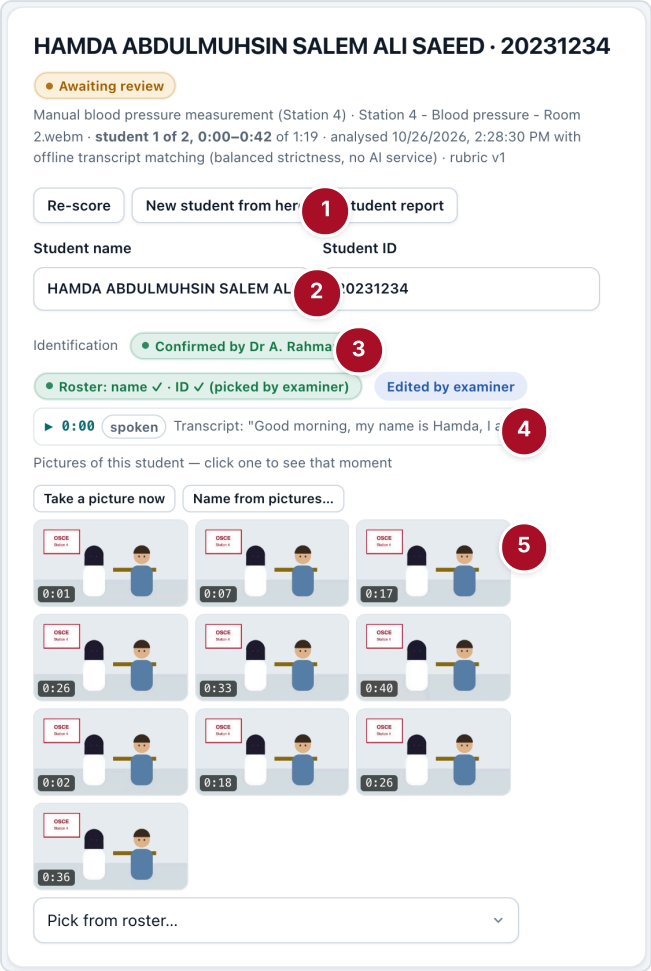


Figure 5.2 — The identity panel.

1. **New student from here**
   Starts a new student at the current point of the video.

2. **Name and ID**
   Typed changes are matched against the roster.

3. **Identification**
   Confirmed by the examiner, or the match confidence.

4. **Evidence**
   ▶ plays the moment the student said their name.

5. **Pictures**
   Click one to jump to that moment.

> **CHECK THIS FIRST**
>
> A score in the wrong student's name is the most serious mistake possible. **Pick from roster...** below the pictures replaces the student in one step.

The transcript

Every line shows its time and who said it. Click a line, or any word, to play from exactly there. The word being spoken is underlined as the video plays.

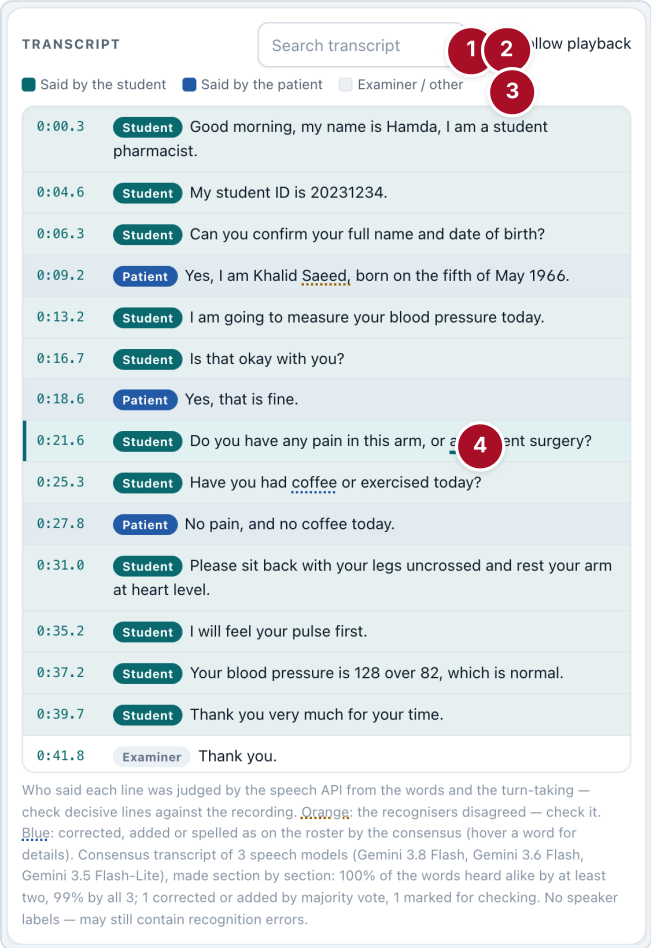


Figure 5.3 — The transcript card, with "recent" underlined at 0:24.

1. **Search**
   Shows only the lines with your words.
2. **Follow playback**
   Keeps the current line in view.
3. **Colour key**
   Teal student, blue patient, grey examiner.
4. **The current word**
   Underlined in step with the video.

YOU SEE	IT MEANS
Teal line	Said by the student. Only these lines earn credit.
Blue line	Said by the patient. Never credited.
Orange word	The recognisers disagreed — listen to it.
Blue word	Corrected or added by the vote.

Confidence of each item

Every item now shows how well the student's own words matched it, as a percentage. The percentage decides the confidence label:

CONFIDENCE	MATCH	WHAT TO DO
Low confidence	under 30%	Little or nothing of the item was heard. Check the video, especially for items that must be seen.
Medium confidence	30% to 60%	Part of the item was heard. Listen to the evidence and decide the level.
High confidence	over 60%	The student's words clearly match the item. A quick check is enough.

Filtering by confidence

Above the items, the filters now include **High confidence**, **Medium confidence** and **Low confidence**, each with its number of items, as well as **Not high confidence** for Medium and Low together.

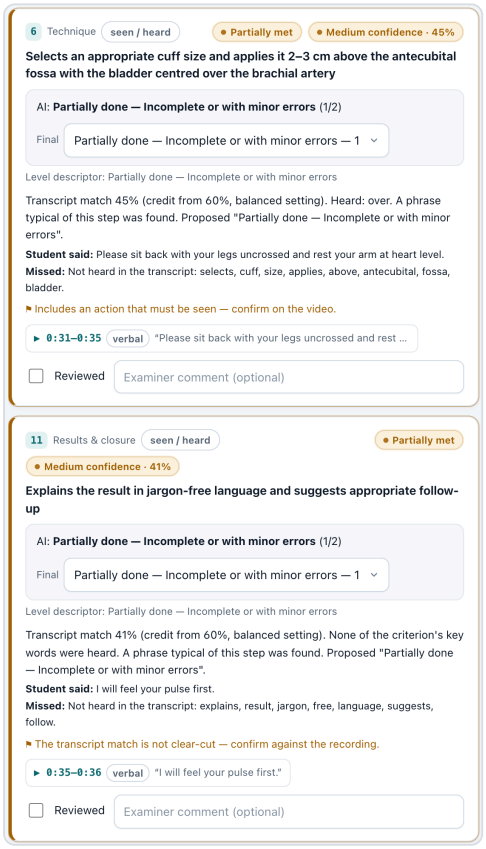


Figure 5.5 — Medium confidence chosen: only the items matched between 30% and 60% are shown.

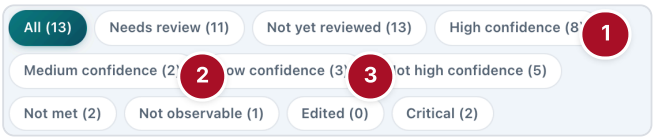


Figure 5.4 — The filters, with the new confidence filters.

START WITH LOW

Work through **Low confidence** first, then **Medium**: these are the items where the transcript says least, and where the video matters most.

Checking an item

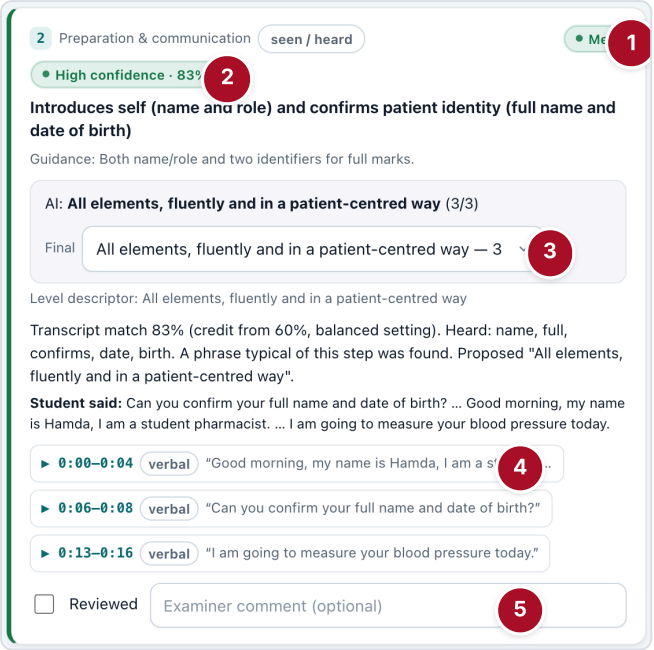


Figure 5.6 — An item matched at 83% — High confidence — with the student's own words as evidence.

1. **Status**
   Met, partially met, not met or not observable.
2. **Confidence and percentage**
   How well the student's words matched.
3. **Final**
   Choose a different level to change the score.
4. **Evidence**
   ▶ plays the exact clip.
5. **Reviewed and comment**
   Tick when checked; the comment appears on the report.

Changing a score

Choose the right level in **Final**. The item is marked reviewed and edited, the totals change at once, and the change is kept in the item's history.

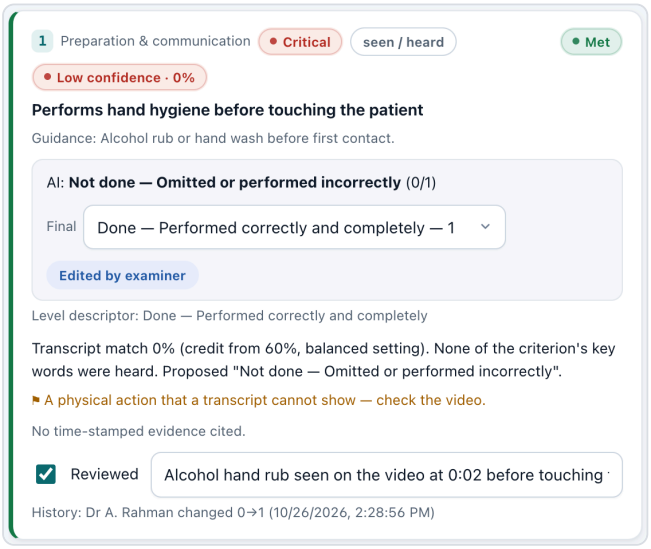


Figure 5.7 — Hand hygiene cannot be heard, only seen. After watching the video the examiner changed it to Done and wrote why.

Overall judgement and feedback



Figure 5.8 — The feedback after the examiner's change: the summary already names it.

1. **Global rating**
   The examiner's own 1–5 rating.
2. **Summary**
   Items met, total and outcome, every score changed, the rating and comments.
3. **Strengths**
   Items met, with the examiner's comments.
4. **Areas for improvement**
   Items partly met or not met, with what was missed.
5. **Examiner comments**
   Your overall comments for the student.

FEEDBACK FOLLOWS YOUR REVIEW

Summary, Strengths and Areas for improvement rewrite themselves as you review. A box you edit yourself is kept as written; empty it to make it automatic again.

Signing off

Sign-off is the examiner's statement that the scores are their own professional judgement. It is recorded on the report and locks the scores.

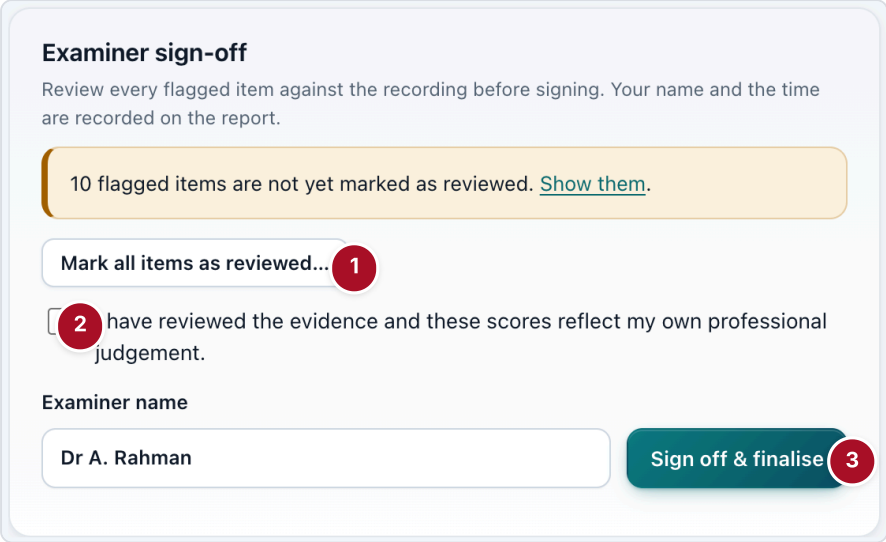


Figure 5.9 — Examiner sign-off.

1. **Mark all items as reviewed**
   Only after checking them.
2. **Confirm**
   Tick the declaration.
3. **Sign off & finalise**
   With your name.

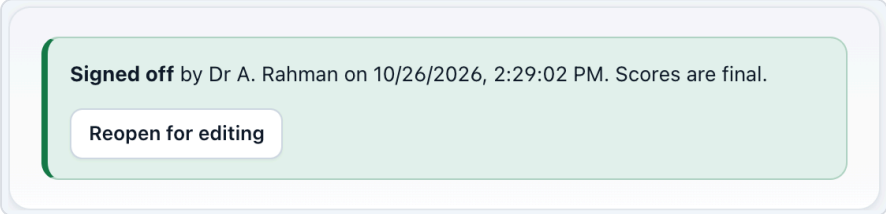


Figure 5.10 — A signed-off student. **Reopen for editing** unlocks the scores.

STATUS	MEANING
Not analysed	Waiting to be analysed.
Awaiting review	Proposed scores ready for an examiner.
Verified	Signed off — the scores are final.
Error	Something failed; the message says what.

CHAPTER SIX

Reports

One report per student and a cohort summary for the whole session. Fill in the report details — now including the exam mark — once, and they are used on every report.

For **coordinators and examiners**

Report details

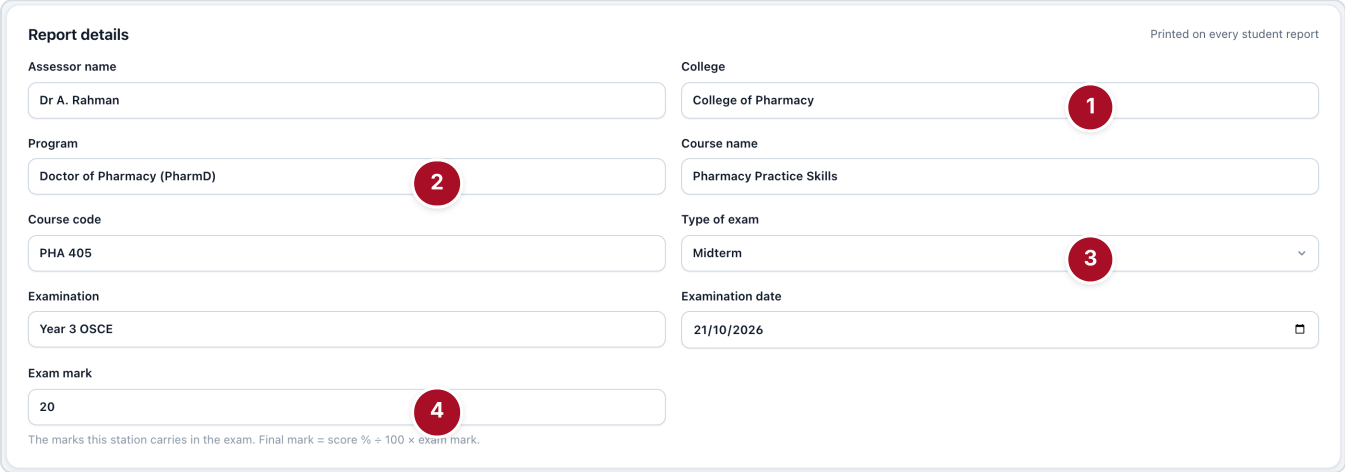


Figure 6.1 — Report details, with the new **Exam mark** field.

FIELD	EXAMPLE	USED FOR
College and Program	College of Pharmacy · Doctor of Pharmacy (PharmD)	The report header
Course name and code	Pharmacy Practice Skills · PHA 405	"Course: … (PHA 405)"
Type of exam	Midterm or End semester	With the examination name and date
Exam mark	20	The final mark of every student

Exam mark and final mark

The **exam mark** is the number of marks this station carries in the exam. Each student's **final mark** is their score percentage applied to it:

> **Final mark = score % ÷ 100 × exam mark**

STUDENT	SCORE	SCORE %	EXAM MARK	FINAL MARK
Student A	18 / 26	69.2%	20	13.85 / 20
Student B	13 / 26	50.0%	20	10 / 20

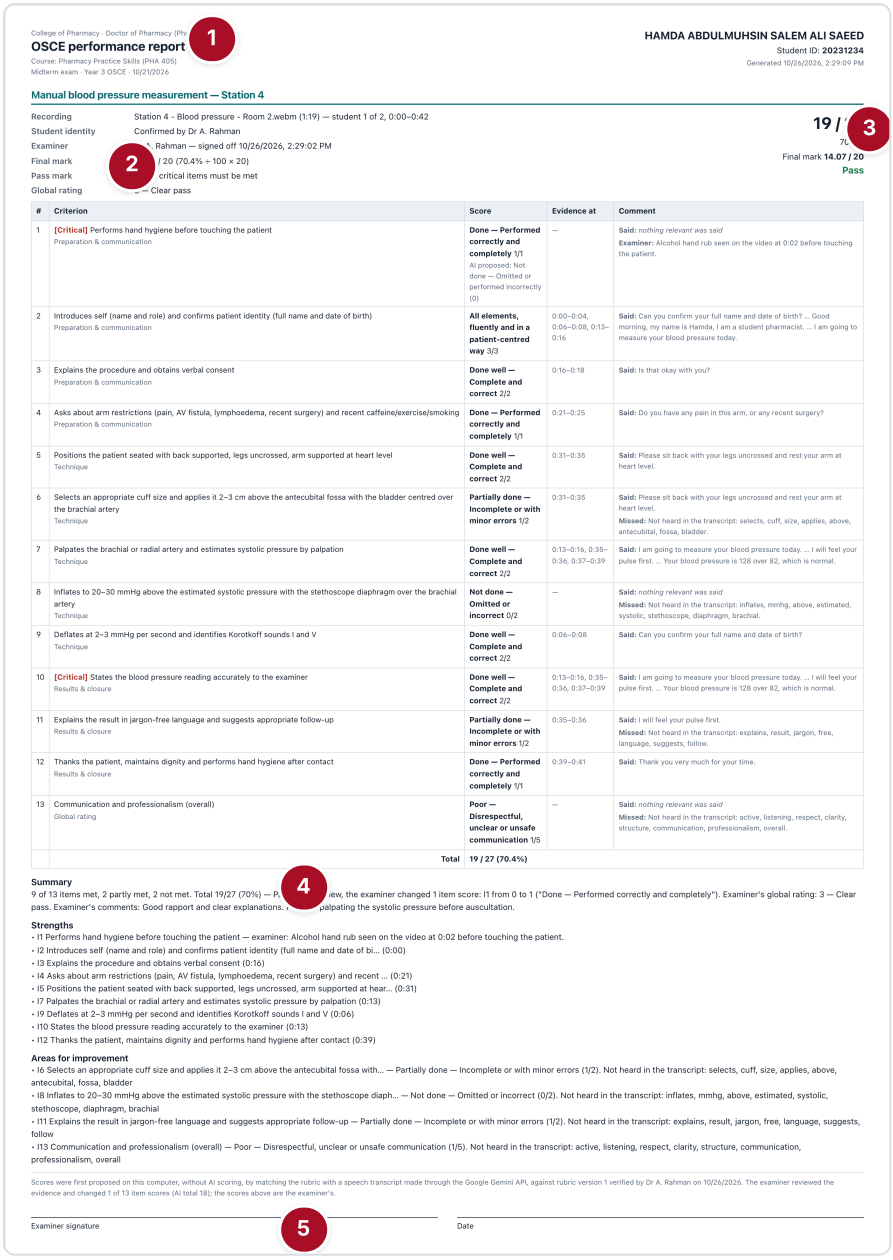
The final mark uses the score **after review**, so it changes when the examiner changes an item. It appears in four places:

WHERE	SHOWN AS
The student report	"Final mark 13.85 / 20" under the score, and in the station details
The overview of a multi-station report	A **Final mark** column
The cohort summary	A **Final mark (/20)** column
Cohort CSV and item-level CSV	**Exam mark** and **Final mark** columns

LEAVE IT EMPTY IF YOU DO NOT NEED IT

Without an exam mark, the final mark is simply not shown, and the reports look as before.

The student report



College of Pharmacy · Doctor of Pharmacy (Pha

OSCE performance report

Course: Pharmacy Practice Skills (PHA 405)
Midterm exam · Year 3 OSCE · 10/21/2026

HAMDA ABDULMUHSIN SALEM ALI SAEED
Student ID: **20231234**
Generated 10/26/2026, 2:29:09 PM

Manual blood pressure measurement — Station 4

Recording	Station 4 - Blood pressure - Room 2.webm (1:19) — student 1 of 2, 0:00–0:42
Student identity	Confirmed by Dr A. Rahman
Examiner	. Rahman — signed off 10/26/2026, 2:29:02 PM
Final mark	/ 20 (70.4% ÷ 100 × 20)
Pass mark	critical items must be met
Global rating	— Clear pass

19 /
7

Final mark **14.07 / 20**
Pass

#	Criterion	Score	Evidence at	Comment
1	[Critical] Performs hand hygiene before touching the patient Preparation & communication	**Done — Performed correctly and completely** 1/1 AI proposed: Not done — Omitted or performed incorrectly (0)	—	**Said:** *nothing relevant was said* **Examiner:** Alcohol hand rub seen on the video at 0:02 before touching the patient.
2	Introduces self (name and role) and confirms patient identity (full name and date of birth) Preparation & communication	**All elements, fluently and in a patient-centred way** 3/3	0:00–0:04, 0:06–0:08, 0:13–0:16	**Said:** Can you confirm your full name and date of birth? … Good morning, my name is Hamda, I am a student pharmacist. … I am going to measure your blood pressure today.
3	Explains the procedure and obtains verbal consent Preparation & communication	**Done well — Complete and correct** 2/2	0:16–0:18	**Said:** Is that okay with you?
4	Asks about arm restrictions (pain, AV fistula, lymphoedema, recent surgery) and recent caffeine/exercise/smoking Preparation & communication	**Done — Performed correctly and completely** 1/1	0:21–0:25	**Said:** Do you have any pain in this arm, or any recent surgery?
5	Positions the patient seated with back supported, legs uncrossed, arm supported at heart level Technique	**Done well — Complete and correct** 2/2	0:31–0:35	**Said:** Please sit back with your legs uncrossed and rest your arm at heart level.
6	Selects an appropriate cuff size and applies it 2–3 cm above the antecubital fossa with the bladder centred over the brachial artery Technique	**Partially done — Incomplete or with minor errors** 1/2	0:31–0:35	**Said:** Please sit back with your legs uncrossed and rest your arm at heart level. **Missed:** Not heard in the transcript: selects, cuff, size, applies, above, antecubital, fossa, bladder.
7	Palpates the brachial or radial artery and estimates systolic pressure by palpation Technique	**Done well — Complete and correct** 2/2	0:13–0:16, 0:35–0:36, 0:37–0:39	**Said:** I am going to measure your blood pressure today. … I will feel your pulse first. … Your blood pressure is 128 over 82, which is normal.
8	Inflates to 20–30 mmHg above the estimated systolic pressure with the stethoscope diaphragm over the brachial artery Technique	**Not done — Omitted or incorrect** 0/2	—	**Said:** *nothing relevant was said* **Missed:** Not heard in the transcript: inflates, mmhg, above, estimated, systolic, stethoscope, diaphragm, brachial.
9	Deflates at 2–3 mmHg per second and identifies Korotkoff sounds I and V Technique	**Done well — Complete and correct** 2/2	0:06–0:08	**Said:** Can you confirm your full name and date of birth?
10	[Critical] States the blood pressure reading accurately to the examiner Results & closure	**Done well — Complete and correct** 2/2	0:13–0:16, 0:35–0:36, 0:37–0:39	**Said:** I am going to measure your blood pressure today. … I will feel your pulse first. … Your blood pressure is 128 over 82, which is normal.
11	Explains the result in jargon-free language and suggests appropriate follow-up Results & closure	**Partially done — Incomplete or with minor errors** 1/2	0:35–0:36	**Said:** I will feel your pulse first. **Missed:** Not heard in the transcript: explains, result, jargon, free, language, suggests, follow.
12	Thanks the patient, maintains dignity and performs hand hygiene after contact Results & closure	**Done — Performed correctly and completely** 1/1	0:39–0:41	**Said:** Thank you very much for your time.
13	Communication and professionalism (overall) Global rating	**Poor — Disrespectful, unclear or unsafe communication** 1/5	—	**Said:** *nothing relevant was said* **Missed:** Not heard in the transcript: active, listening, respect, clarity, structure, communication, professionalism, overall.
		Total	**19 / 27 (70.4%)**	

Summary
9 of 13 items met, 2 partly met, 2 not met. Total 19/27 (70%) — P… …iew, the examiner changed 1 item score: I1 from 0 to 1 ("Done — Performed correctly and completely"). Examiner's global rating: 3 — Clear pass. Examiner's comments: Good rapport and clear explanations. … palpating the systolic pressure before auscultation.

Strengths
- I1 Performs hand hygiene before touching the patient — examiner: Alcohol hand rub seen on the video at 0:02 before touching the patient.
- I2 Introduces self (name and role) and confirms patient identity (full name and date of bi… (0:00)
- I3 Explains the procedure and obtains verbal consent (0:16)
- I4 Asks about arm restrictions (pain, AV fistula, lymphoedema, recent surgery) and recent … (0:21)
- I5 Positions the patient seated with back supported, legs uncrossed, arm supported at hear… (0:31)
- I7 Palpates the brachial or radial artery and estimates systolic pressure by palpation (0:13)
- I9 Deflates at 2–3 mmHg per second and identifies Korotkoff sounds I and V (0:06)
- I10 States the blood pressure reading accurately to the examiner (0:13)
- I12 Thanks the patient, maintains dignity and performs hand hygiene after contact (0:39)

Areas for improvement
- I6 Selects an appropriate cuff size and applies it 2–3 cm above the antecubital fossa with… — Partially done — Incomplete or with minor errors (1/2). Not heard in the transcript: selects, cuff, size, applies, above, antecubital, fossa, bladder
- I8 Inflates to 20–30 mmHg above the estimated systolic pressure with the stethoscope diaph… — Not done — Omitted or incorrect (0/2). Not heard in the transcript: inflates, mmhg, above, estimated, systolic, stethoscope, diaphragm, brachial
- I11 Explains the result in jargon-free language and suggests appropriate follow-up — Partially done — Incomplete or with minor errors (1/2). Not heard in the transcript: explains, result, jargon, free, language, suggests, follow
- I13 Communication and professionalism (overall) — Poor — Disrespectful, unclear or unsafe communication (1/5). Not heard in the transcript: active, listening, respect, clarity, structure, communication, professionalism, overall

Scores were first proposed on this computer, without AI scoring, by matching the rubric with a speech transcript made through the Google Gemini API, against rubric version 1 verified by Dr A. Rahman on 10/26/2026. The examiner reviewed the evidence and changed 1 of 13 item scores (AI total 18); the scores above are the examiner's.

Examiner signature

Date

Figure 6.2 — A student report with the final mark beside the score and in the station details.

1 Header
College, program, course, exam type and date.

2 Final mark
Score % ÷ 100 × exam mark.

3 Score
With the final mark beneath it.

4 Feedback
Reflects the review.

5 Signature
Examiner and date.

The cohort summary

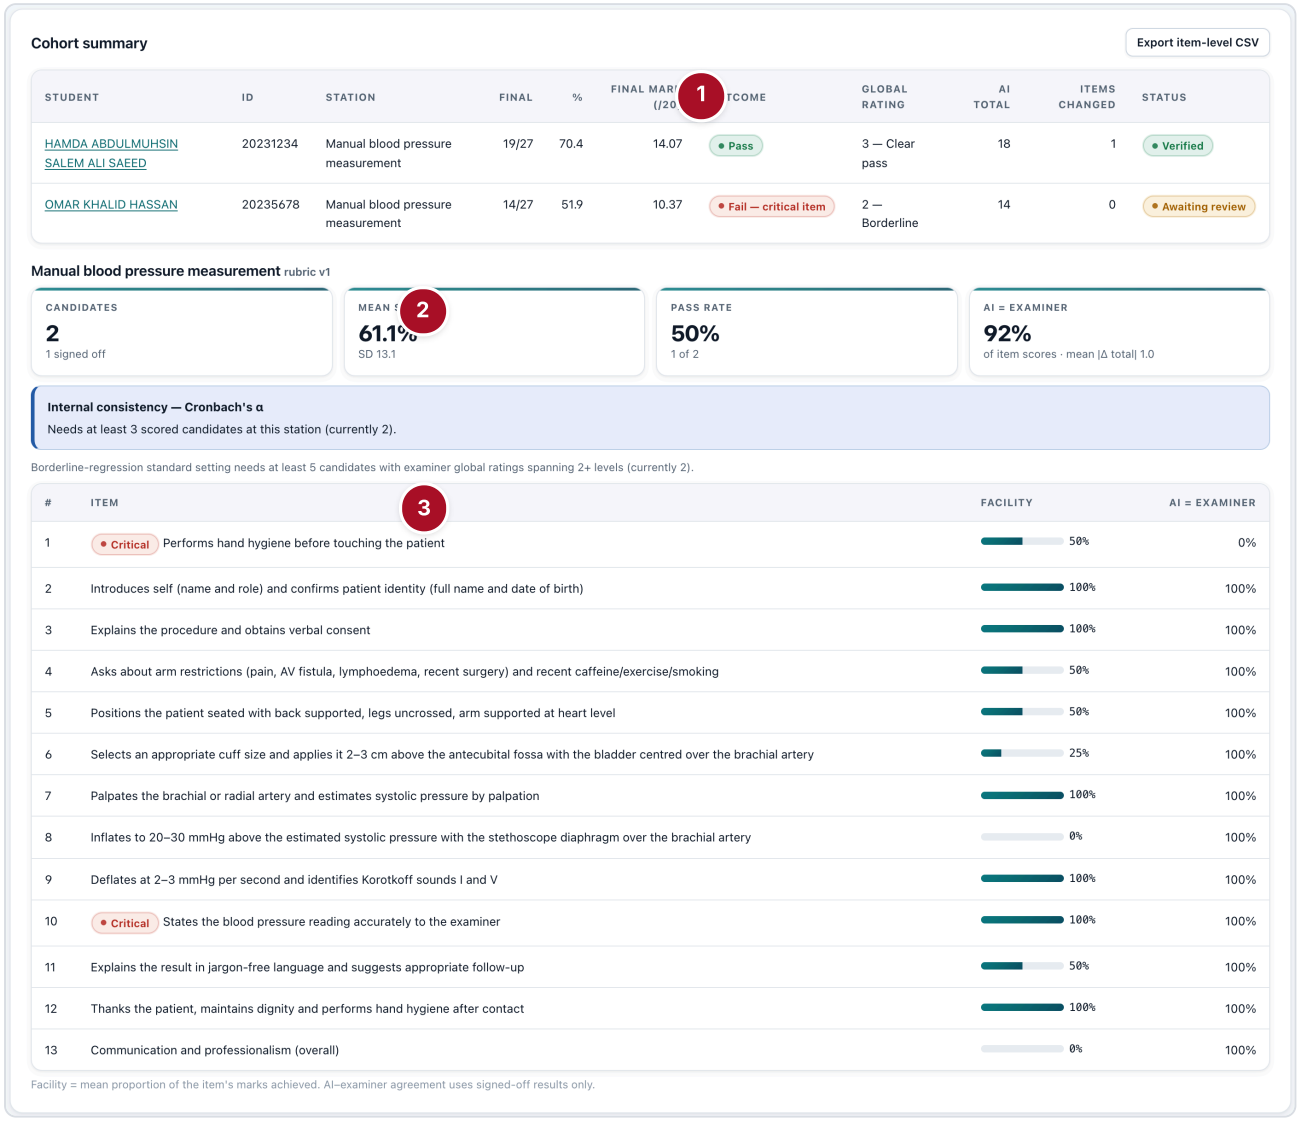


Figure 6.3 — The cohort summary, with the Final mark column.

1 **Final mark**
Each student's final mark out of the exam mark.

2 **Station statistics**
Candidates, mean, pass rate, AI–examiner agreement.

3 **Items**
The facility of each item.

Exporting the results

BUTTON	WHAT YOU GET
Export cohort CSV	One row per student and station: totals, outcome, global rating, examiner, sign-off time, **exam mark and final mark**.
Export item-level CSV	One row per student: name, ID, total after review, %, **final mark**, then the final score of every item.
Print all reports	Every student report, ready to save as one PDF.
Export full backup (.json)	The whole session, to import later and continue.
Export grades for moderation (.json)	Each student's score and %, for another grader to import (page 34).

The item-level CSV

Student name	Student ID	Total	Max	%	Final mark (/20)	Item 1	Item 2	...
HAMDA ABDULMUHSIN SALEM ALI SAEED	20231234	18	26	69.2	13.85	1	2	...
OMAR KHALID HASSAN	20235678	13	26	50.0	10	0	2	...

Keeping your work

Nothing is kept after the browser tab closes. Before closing, export what you need — or press **Export full backup (.json)** and **Import backup** later to continue. Recordings are not inside the backup; re-attach them after importing.

Moderation

When a station is marked by two examiners, the application compares their results and agrees one result per student. Moderation is **optional**: without a second grader's file, the reports stay exactly as they are.

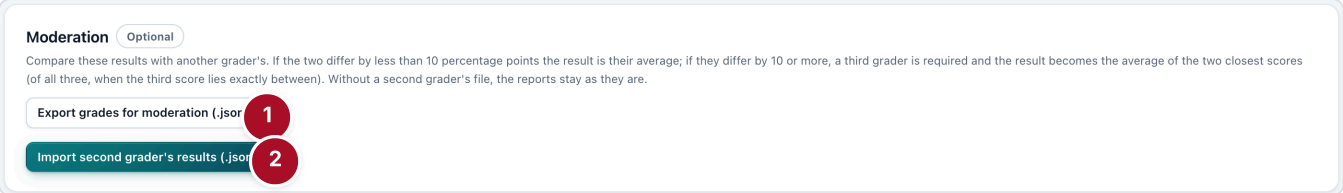


Figure 6.4 — The Moderation card at the top of the cohort summary, before a second grader's results are imported.

1. **Export grades for moderation**
   Each grader exports their own results as a small .json file and sends it to the moderator.

2. **Import second grader's results**
   The moderator imports the other grader's file. A full backup of their session works too.

How the result is agreed

THE TWO GRADERS DIFFER BY	THE MODERATED RESULT IS	SHOWN AS
less than 10 percentage points	the average of the two	Agreed — average
10 percentage points or more	not decided yet — a third grader is required	Third grader required
10 or more, with a third grader	the average of the two closest scores of the three (of all three when the third lies exactly between)	Moderated by third grader

Students are matched by student ID, then by name, and by station. The difference is |grader 1 % − grader 2 %|; the moderated score is the moderated % of the station maximum, and the final mark is the moderated % of the exam mark.

When a third grader is needed

After the second grader's file is imported, the card lists every student with both results and the difference. Differences of 10 points or more are shown in red, and the card asks for a third grader.

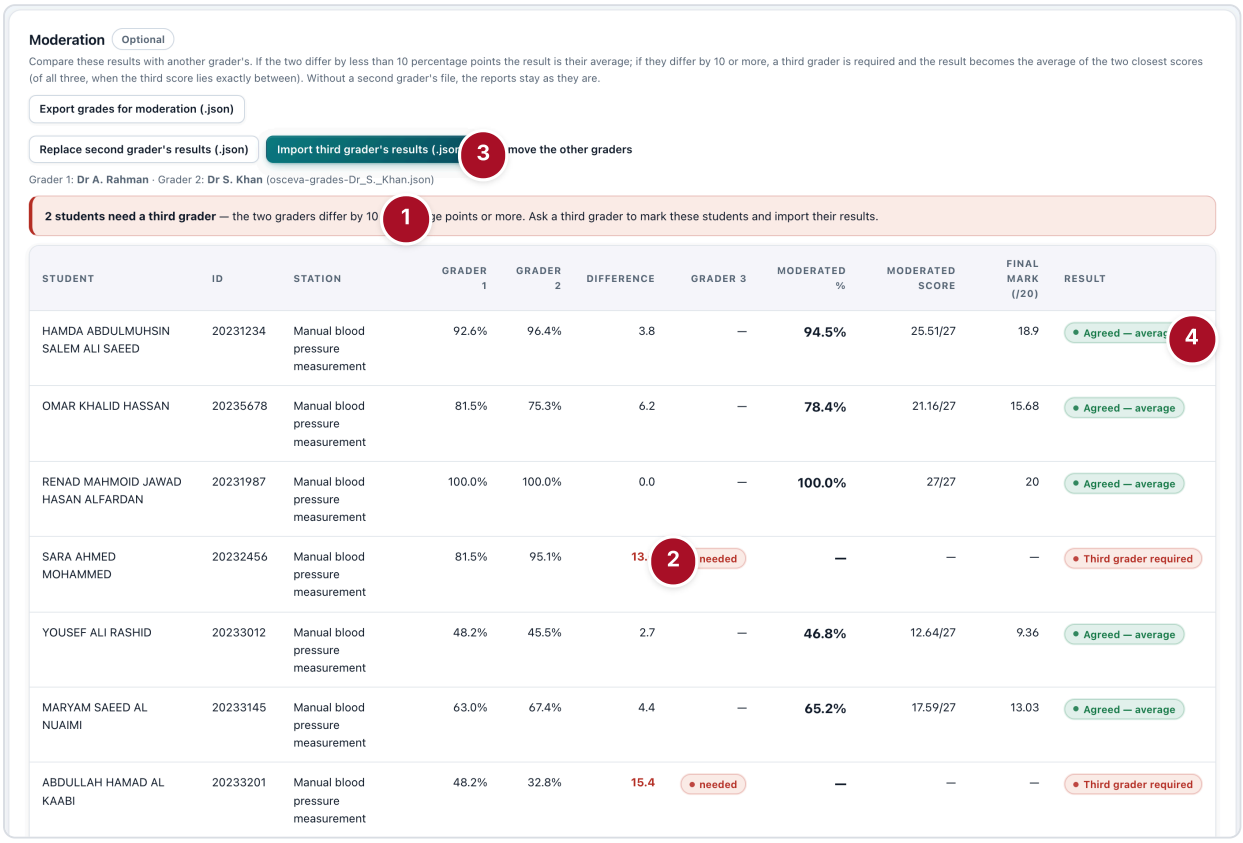


Figure 6.5 — Two students — 13.6 and 15.4 points apart — need a third grader; the others are averaged. The first seven students are shown.

1 **Third grader required**
How many students need a third grader.

2 **The difference**
In red when 10 points or more.

3 **Import third grader's results**
The third grader needs to mark only those students.

4 **Agreed — average**
Under 10 points apart: the average is used.

Importing the third grader

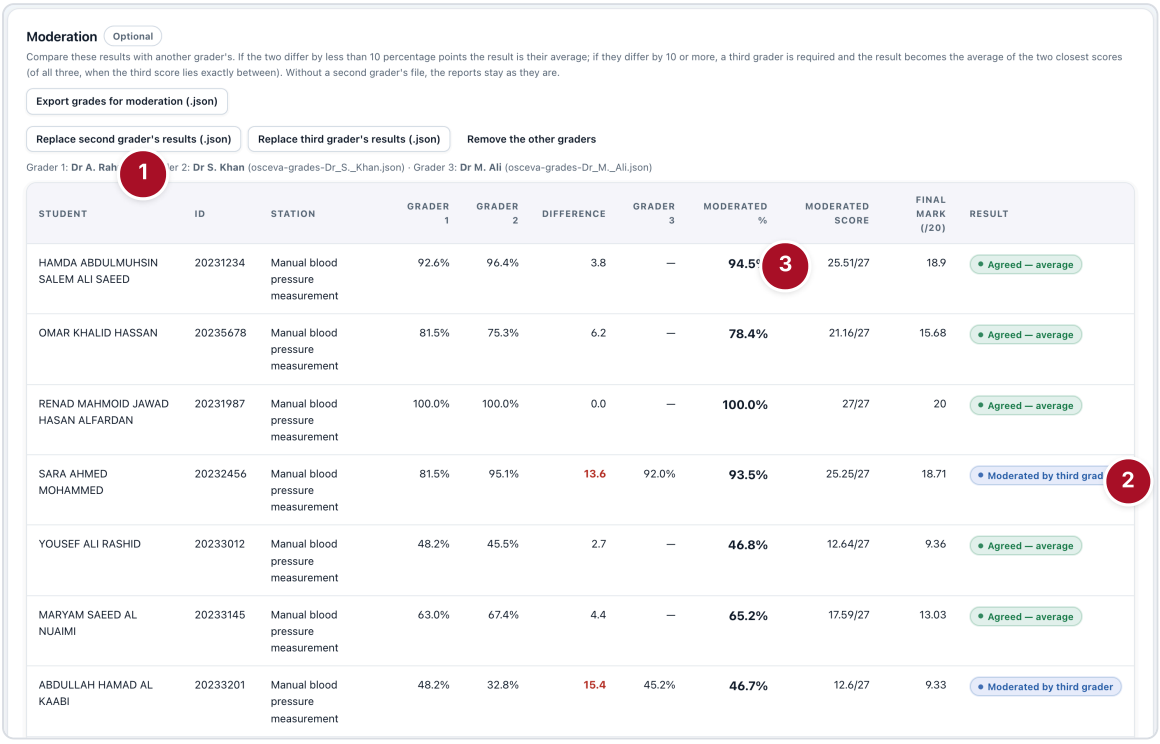


Figure 6.6 — With the third grader's results, every student has a moderated result and final mark (first seven students shown).

1 **The graders**
The name of each grader and the file their results came from.

2 **Moderated by third grader**
The average of the two closest of the three scores.

3 **Moderated %**
The result used for the student.

On the student report

Every report of a moderated student shows the moderation, with each grader's score:

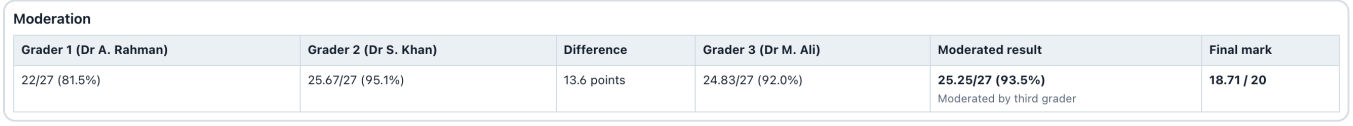


Figure 6.7 — Grader 1 81.5%, grader 2 95.1%, grader 3 92.0%: the two closest (95.1% and 92.0%) are averaged to 93.5%.

REMOVE THE OTHER GRADERS

Remove the other graders returns the reports to this session's own results. Re-importing a file replaces the earlier one.

Cronbach's alpha

The cohort summary now shows, for each station, **Cronbach's alpha** — how consistently the items measure the same performance — with a short conclusion written for you. It needs at least three students.

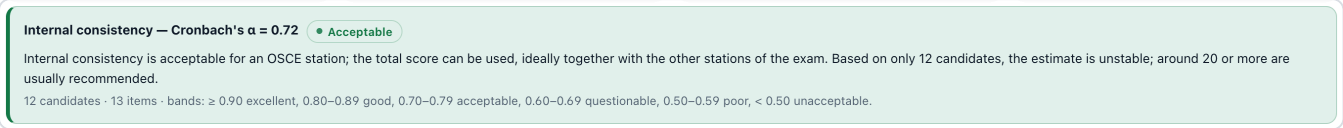


Figure 6.8 — Cronbach's α = 0.72 for 12 candidates and 13 items: acceptable, with a note that 12 candidates is a small cohort.

A	BAND	THE CONCLUSION SAYS
0.90 or more	Excellent	Highly reliable total score. Above 0.95, some items may repeat each other.
0.80 – 0.89	Good	Reliable enough for summative decisions.
0.70 – 0.79	Acceptable	Acceptable for an OSCE station; best used with the other stations.
0.60 – 0.69	Questionable	Use with caution; review items with very high or very low facility.
0.50 – 0.59	Poor	Not reliable on its own; review the checklist.
under 0.50	Unacceptable	Not reliable; a negative value often means a scoring or rubric error.

SMALL COHORTS

With fewer than about 20 candidates the estimate is unstable, and the conclusion says so. Alpha is calculated from the final item scores — after review.

CHAPTER SEVEN

Reference

What to do when something does not go as expected, the labels you will see, and good practice.

For everyone

When something goes wrong

IF YOU SEE	THEN
Google AI Studio says it is not available	You are signed in with a university account. Sign in with a personal Gmail address.
The Gemini API key is not valid	Copy the whole key again (it starts AIza) and press Test connection.
The daily Gemini quota is used up	The next model takes over by itself. If all are used up, wait for the reset or enable billing.
Your OpenAI account has no credit	Add credit under Billing on the OpenAI Platform.
Load the class roster first	Load it on the Recordings step.
The rubric must be verified first	Complete the verification box.
Some students could not be re-scored	Their recording is not loaded. Re-attach it, then choose Re-score in its ⋯ menu.
Not recognised	No roster name or ID was heard. Choose the student from the pictures.
No final mark appears	Enter the Exam mark in Report details.
This media format cannot be decoded	Use Chrome or Edge, or export the video as MP4.

Vertex AI and moderation

IF YOU SEE	THEN
The Vertex AI access token is not valid or has expired	Tokens last one hour. Paste a new one, or use an API key.
An access token needs the Google Cloud project ID	Enter the project ID in the Vertex AI steps of Setup.
Third grader required	Ask a third grader to mark the students listed and import their results (Reports → Moderation).
Students have no result in the second grader's file	Their ID or name, or the station, differs between the two files. Check the other grader's session.

Labels you will see

LABEL	MEANING
Met · Partially met · Not met · Not observable	The proposed or final level of an item.
High confidence · 83%	The student's words matched the item by more than 60%.
Medium confidence · 45%	Matched by 30% to 60%.
Low confidence · 18%	Matched by less than 30%.
Critical	Must be met to pass, when the critical-item rule is on.
Edited by examiner	Changed by hand during review.

Privacy at a glance

SENT	NEVER SENT
The sound of each recording, in pieces	The video picture
The transcript text, once, to mark who is speaking	The class roster — names and IDs
	The rubric, scores, comments and reports
	File names, file dates, the page address, cookies and other metadata

Good practice

1. **Personal email for keys**
   University accounts are refused by Google AI Studio.
2. **Say the name and ID**
   Ask students to state both at the start.
3. **Review Low confidence first**
   Then Medium, then the rest.
4. **Moderate high-stakes stations**
   A second grader, and a third when they differ by 10 points or more.
5. **Export before closing**
   The session is not kept.

GMU OSCEva Assessor · College of Pharmacy · Gulf Medical University

API transcription edition · User manual, third edition · October 2026