Digitising the Murten Panorama

Scanning details

- Target resolution: 1000dpi (40 pixels per mm)
- Painting is 10m high by 100m long
- Stored on 3 rolls

- Photographed with PhaseOne iXH (150 MPixels)
- 72mm lens
- CaptureOne DH software
- Lighting with Ajurat D8 area light and employing cross polarisation
- Custom automation control software
- Custom scanning rig

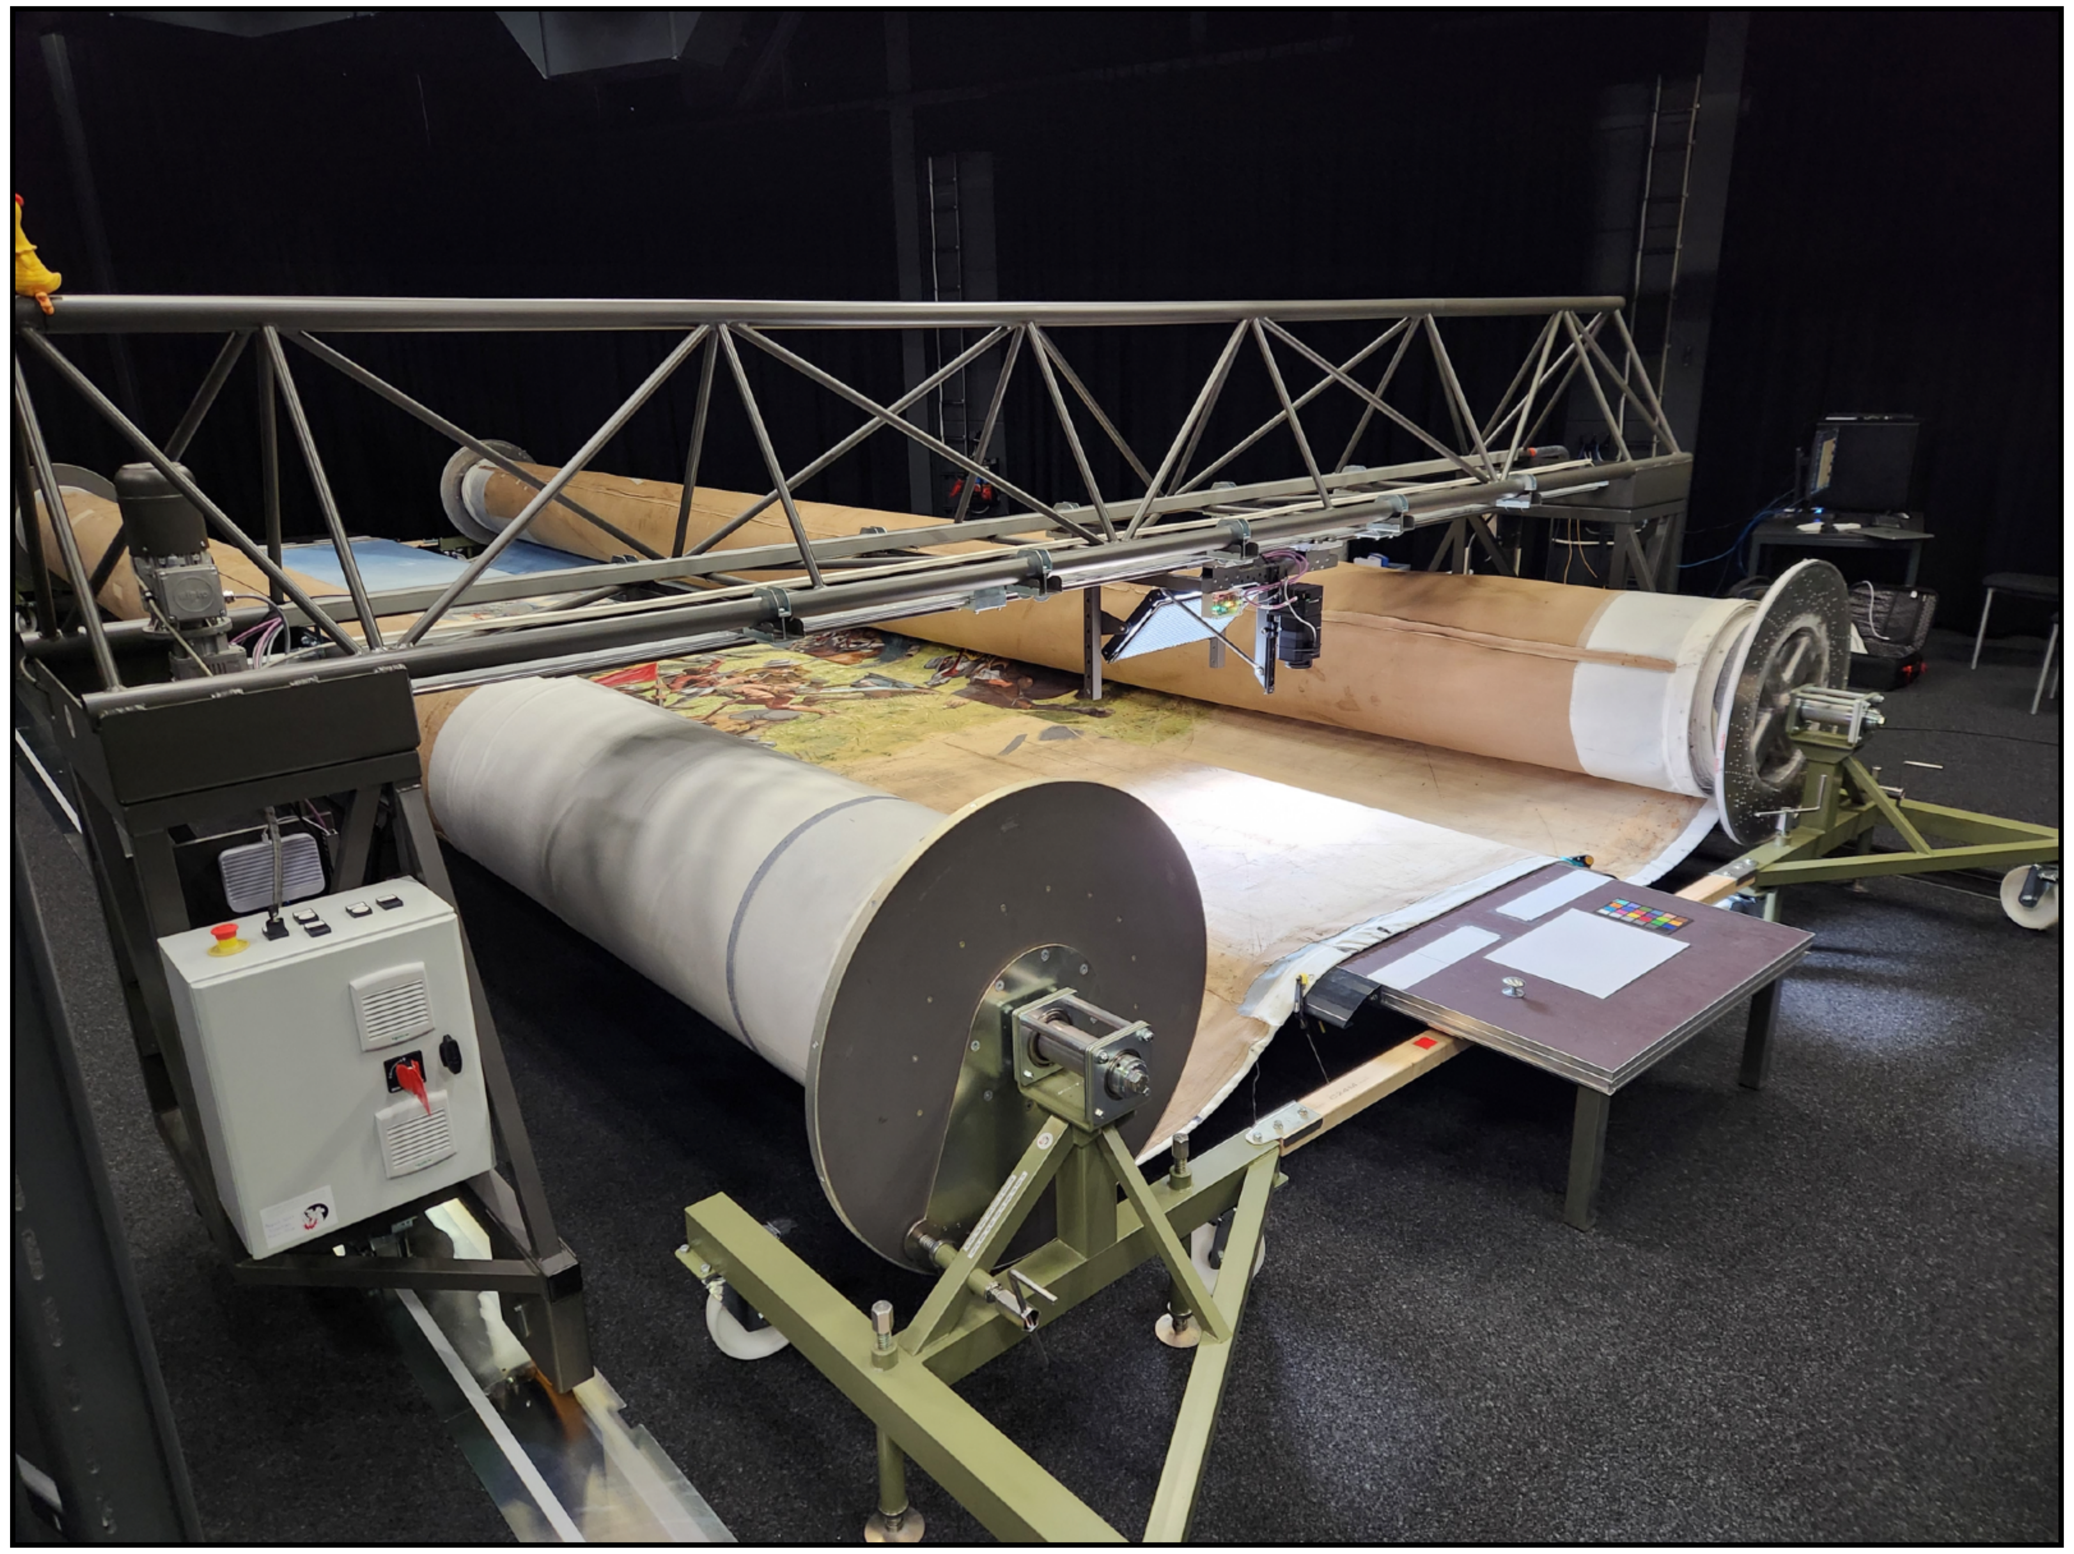

Scanning rig

Achieving the target resolution

- One cannot buy an arbitrary high resolution camera sensor.
- Standard solution is to use multiple photographs and stitch/blend the results together.
- Four general options
  1. Area scan, radial
  2. Area scan, linear
  3. Line scan, radial
  4. Line scan, linear
- Radial scans (1) and (3) can better handle 3D objects.
- For the Murten digitisation the only practical option is (2), linear area scan.
  Greatest risk for linear area scan are 3D elements, eg: brush strokes, causing stitching errors.

Radial vs linear: parallax issue

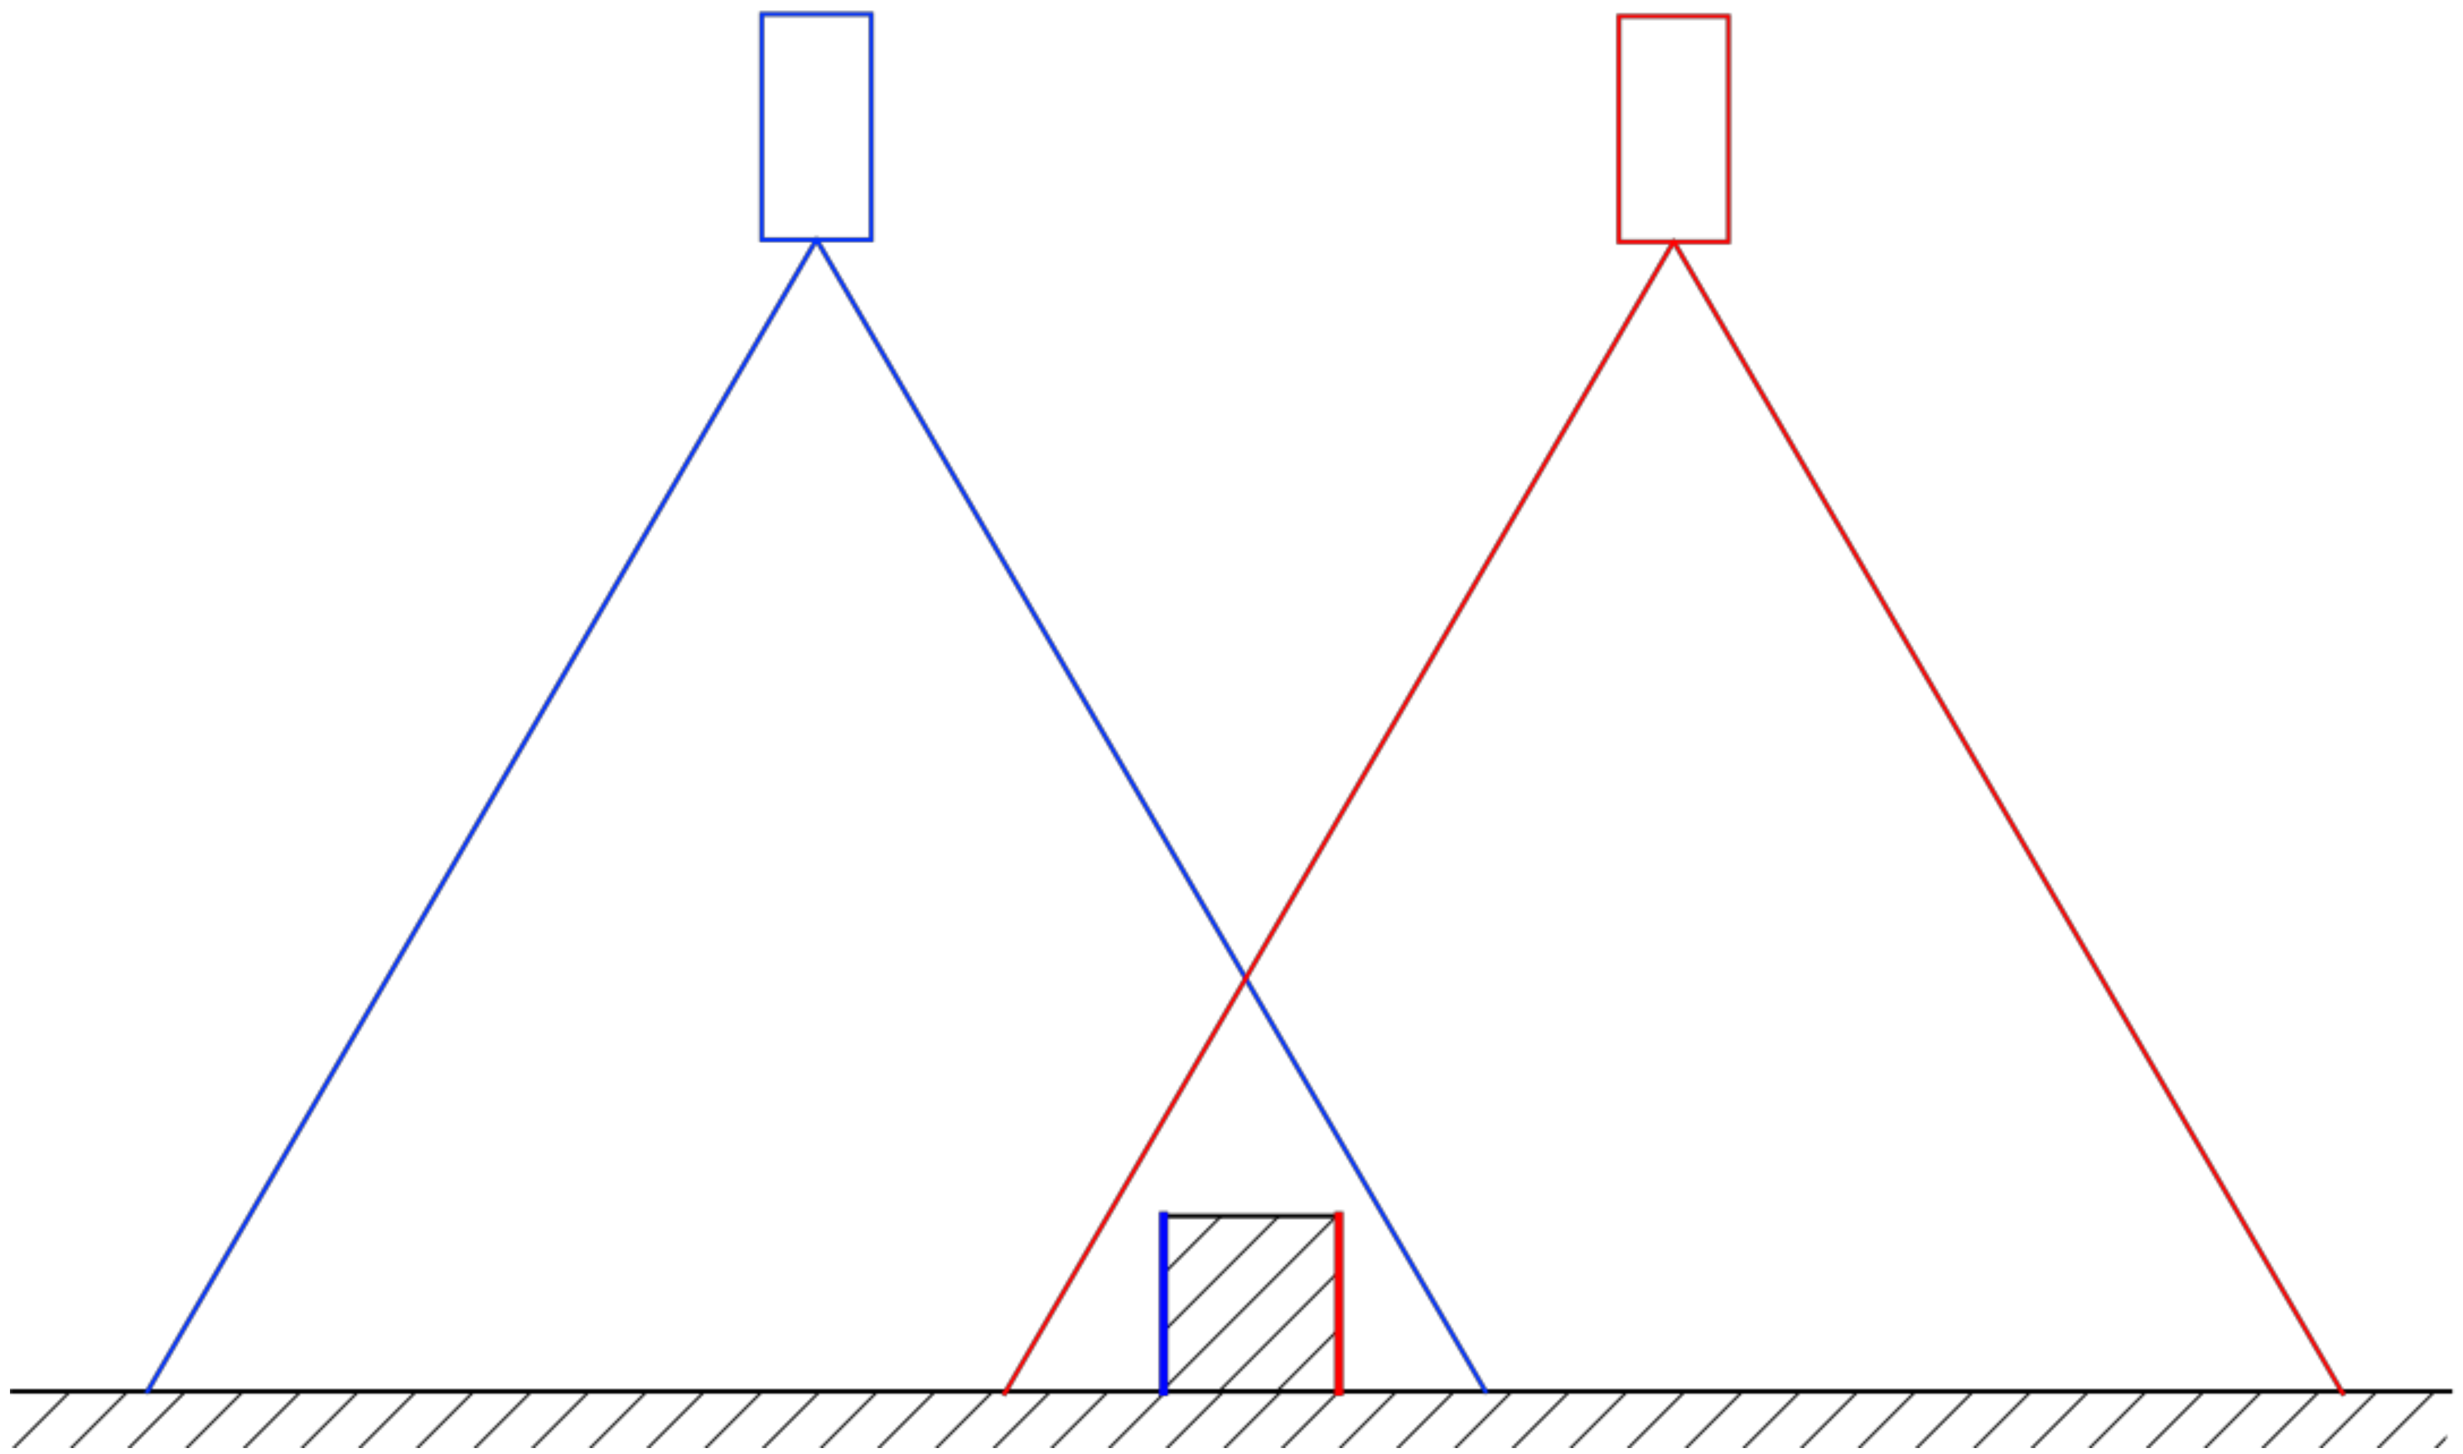

Linear

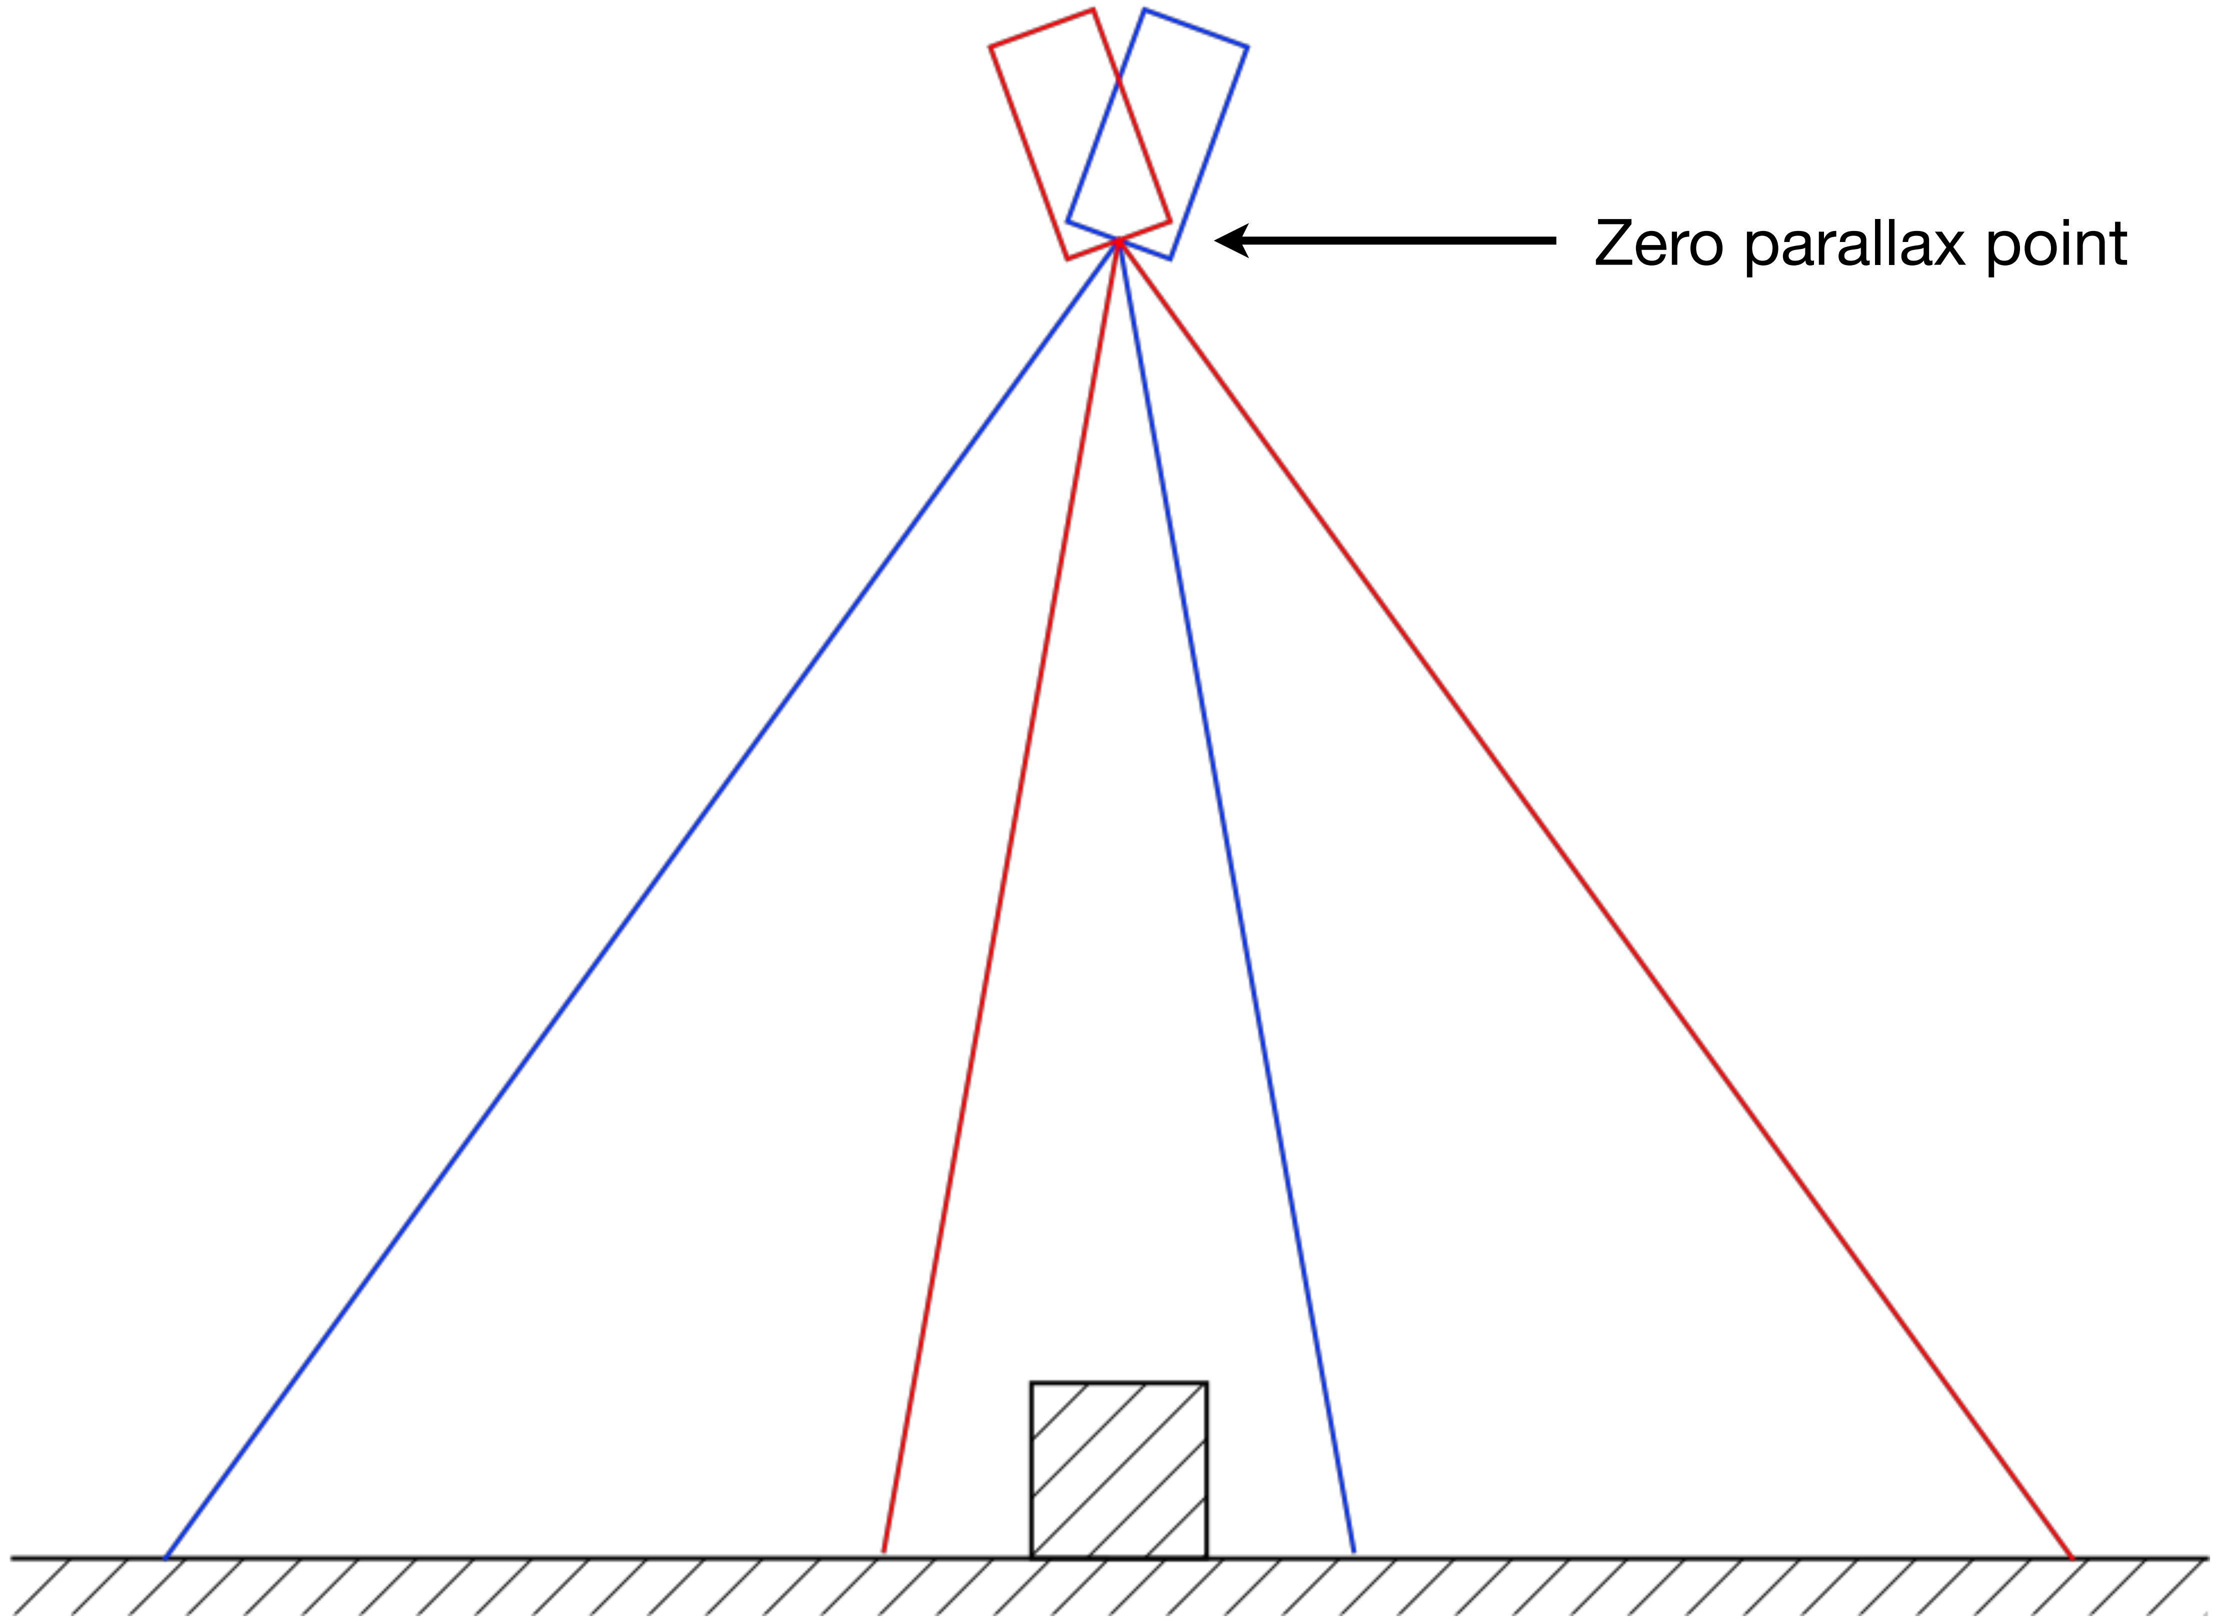

Radial

Area scan, radial: Panorama photography

- Camera rotates about a single position.


Deep Fakes
Lighten Up!
On Biology and Time
EPFL Pavilions


Lighten Up!
On Biology and Time
Uwe Rieger
Yinan Liu
Cyber Physical
Architecture in Real Time
22.9.
–16.6.
EPFL Pavilions
Machines Thinking Musically
20.9.
–29.6.
EPFL Pavilions

Line scanning, linear or radial

- Camera moves (along a line or around a point) and a strip of pixels extracted at each position.
- Method used by most mobile phones when capturing a wide angle panorama.
- True line scanner only has a sensor with just 1 (or a few) pixels wide.

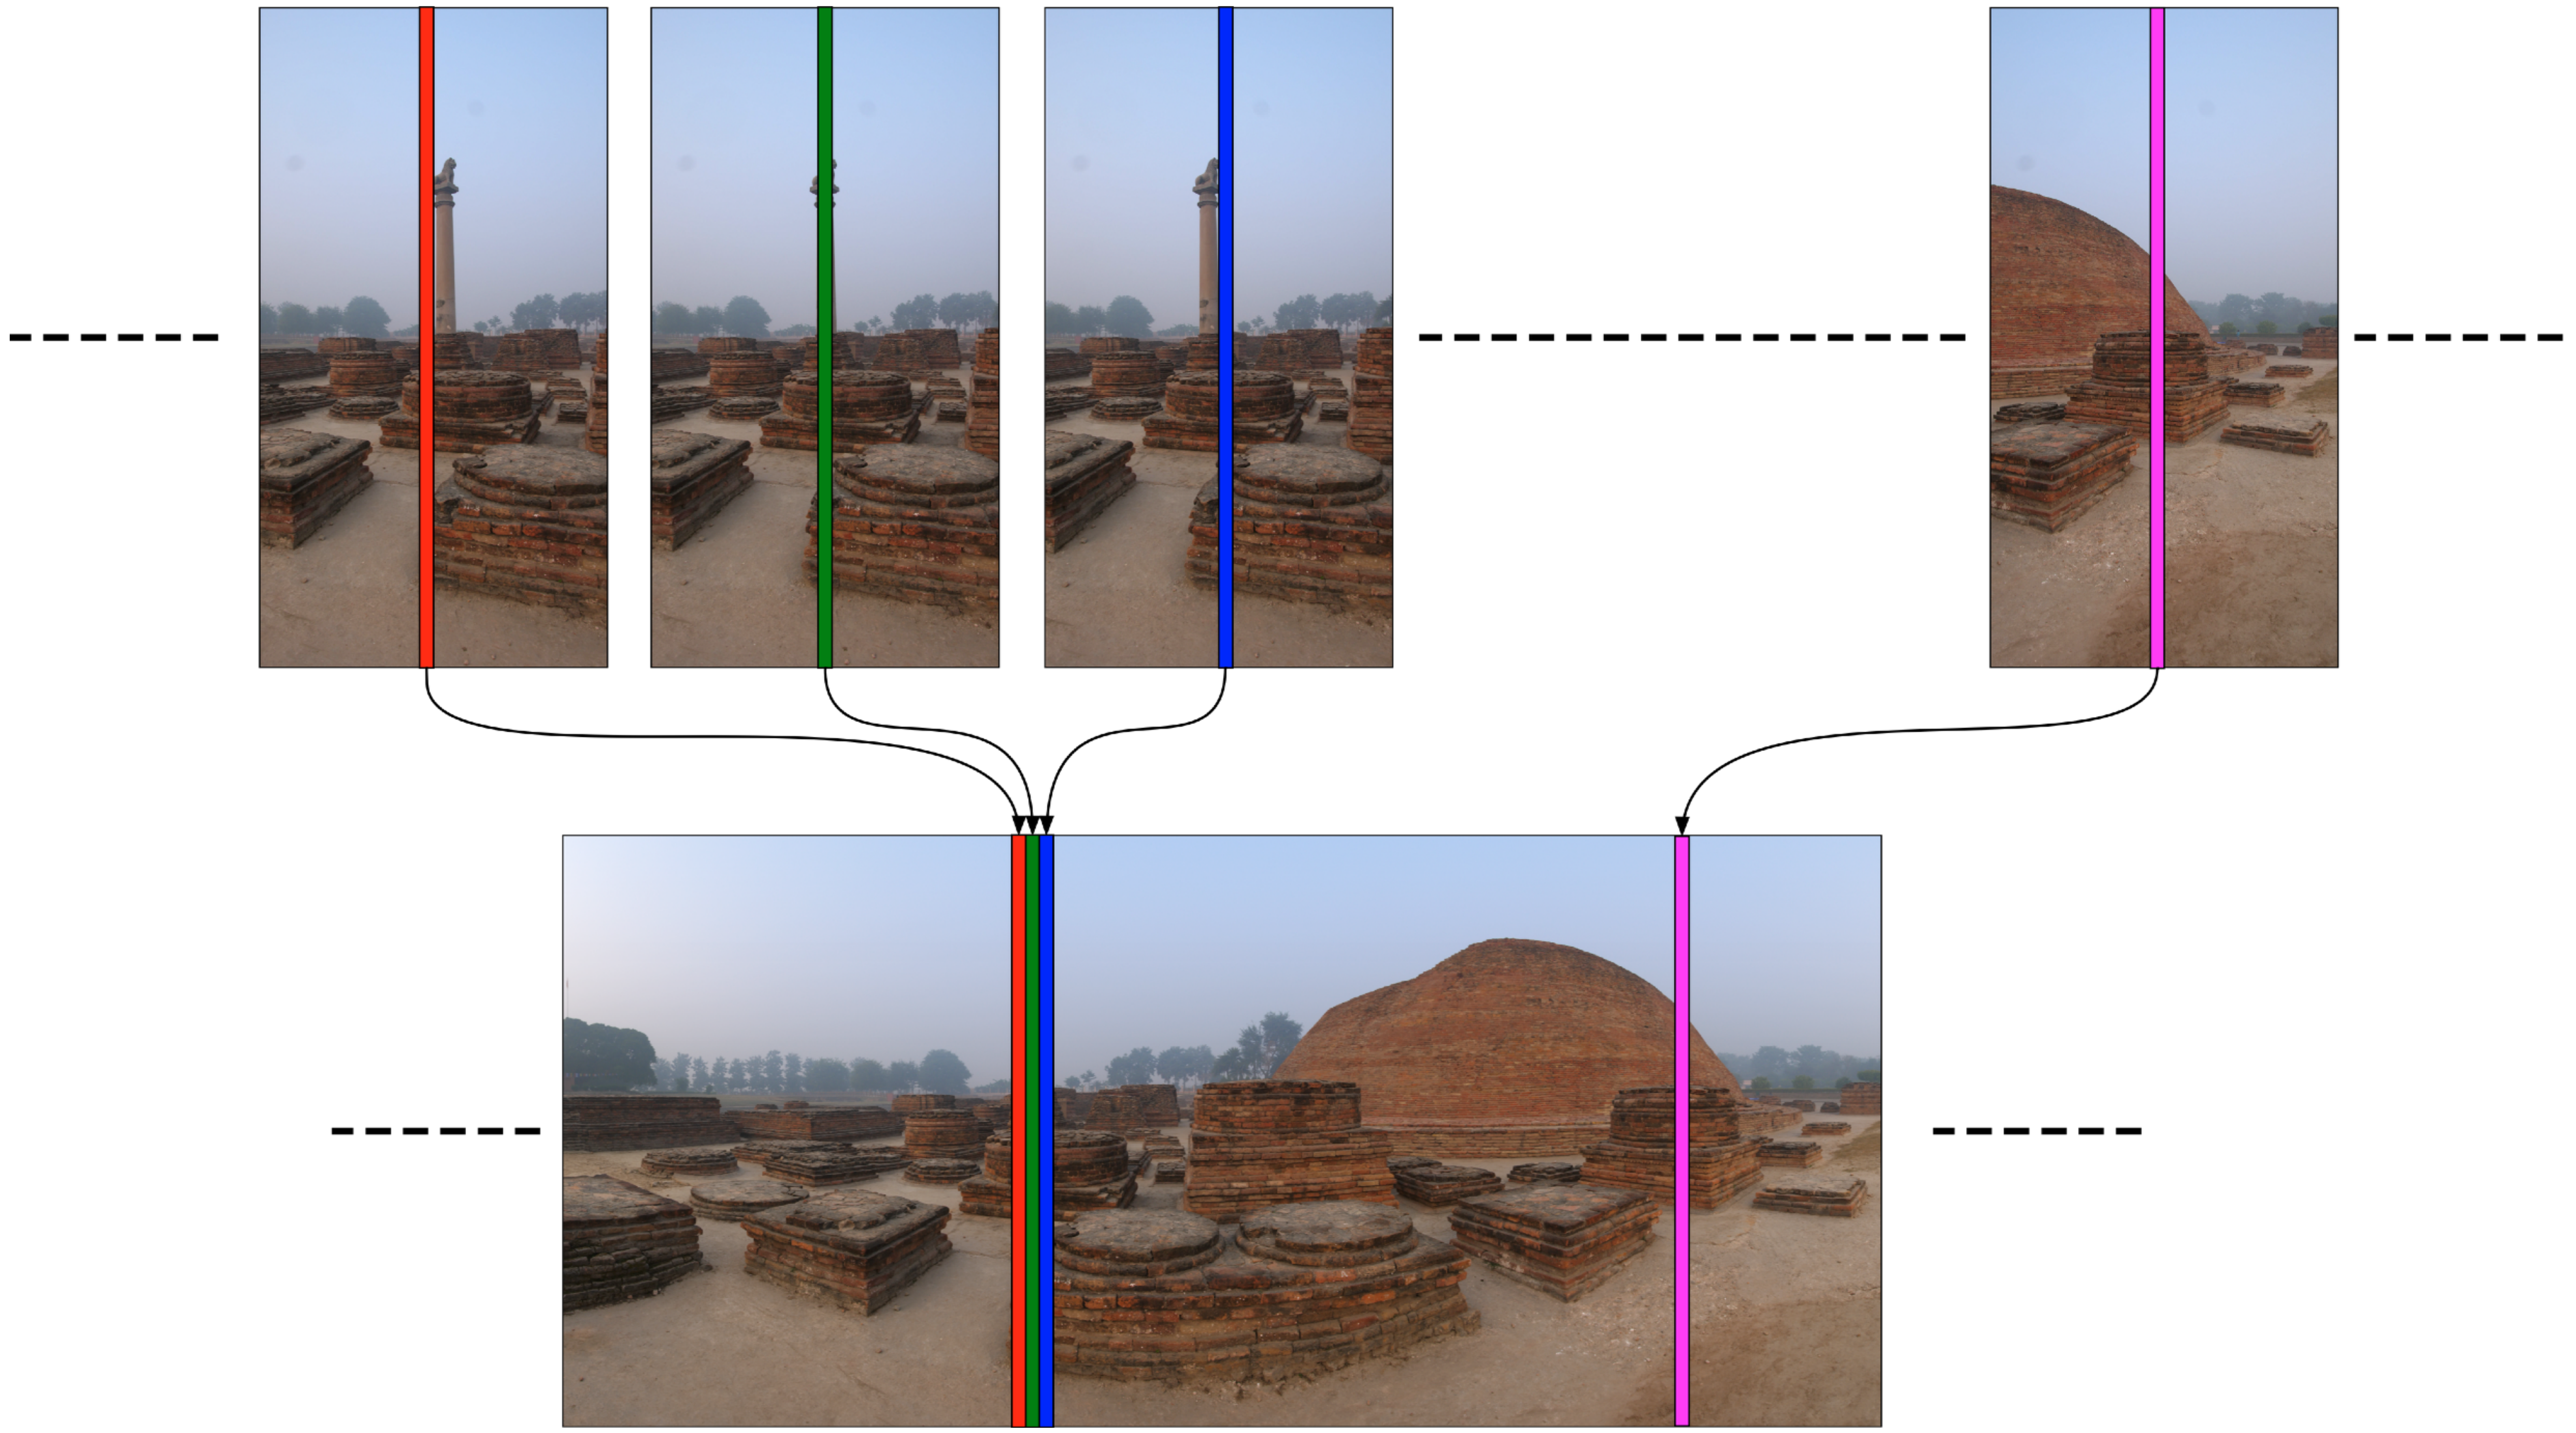


LUMIX
GH5

Area scan, linear example

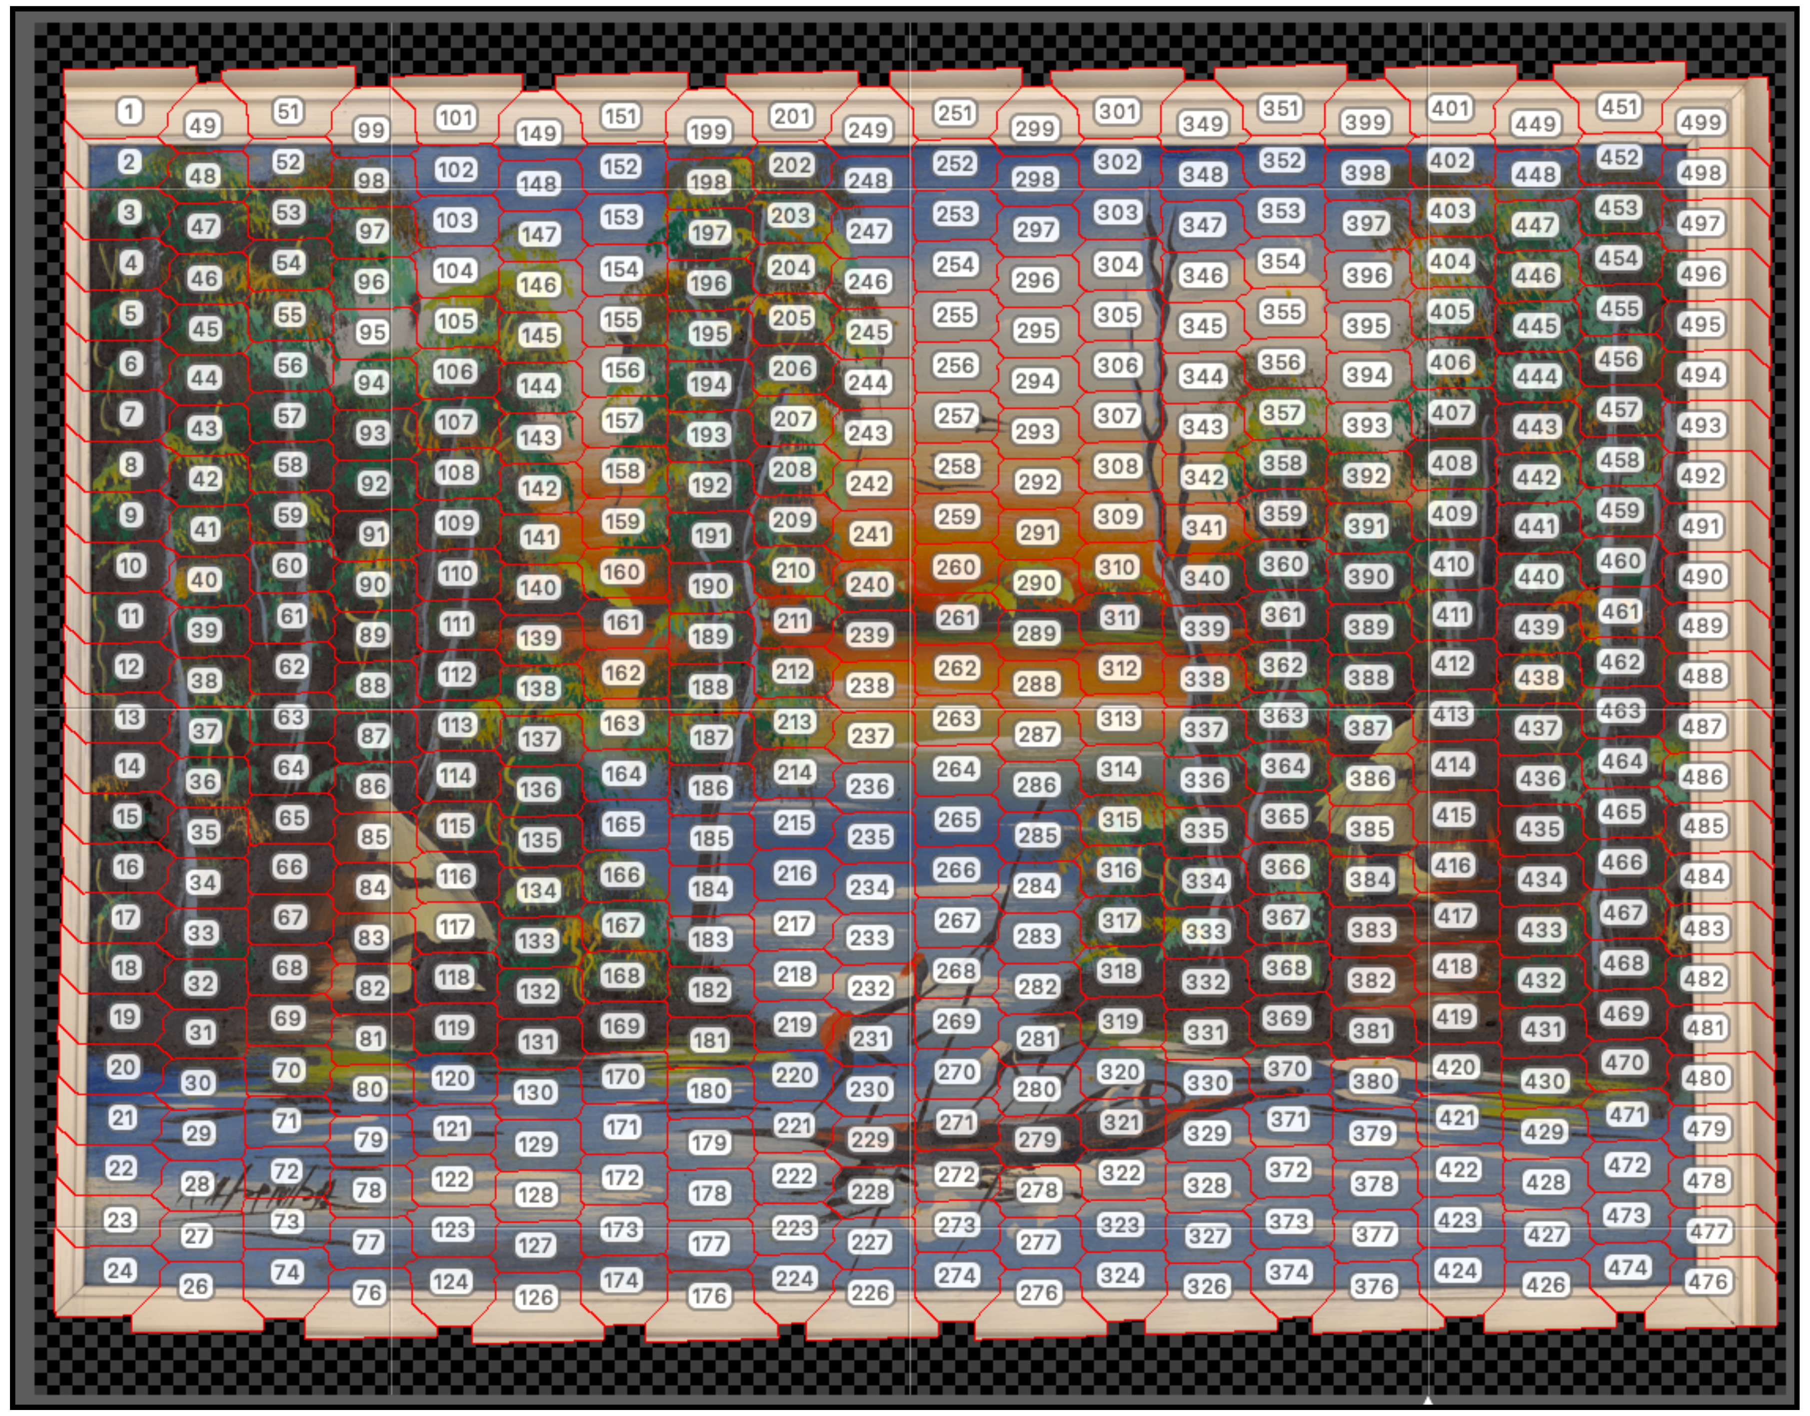

Area scan specifications

- If one knows the target dots per inch, 1000 dpi = 40 pixels per mm

- And one knows the resolution of the sensor, PhaseOne ixH = 14204 x 10652 pixels

- Then one can calculate the width (and height) of each image on the canvas
  14202/40 x 10652/40 = 355 x 266mm

- If one has a lens with a known sensor size and aperture, then one can work out how high the lens should be above the canvas. Sensor width 54 mm, aperture = 72mm then by similar triangles the lens height above the canvas is 355 x 72 / 54 = 473 mm

- Finally one needs to choose the degree of image overlap for the stitch/blend process.
  30% is a semi-standard.

- For the Murten used 30% overlap in one direction and 50% in the other direction.

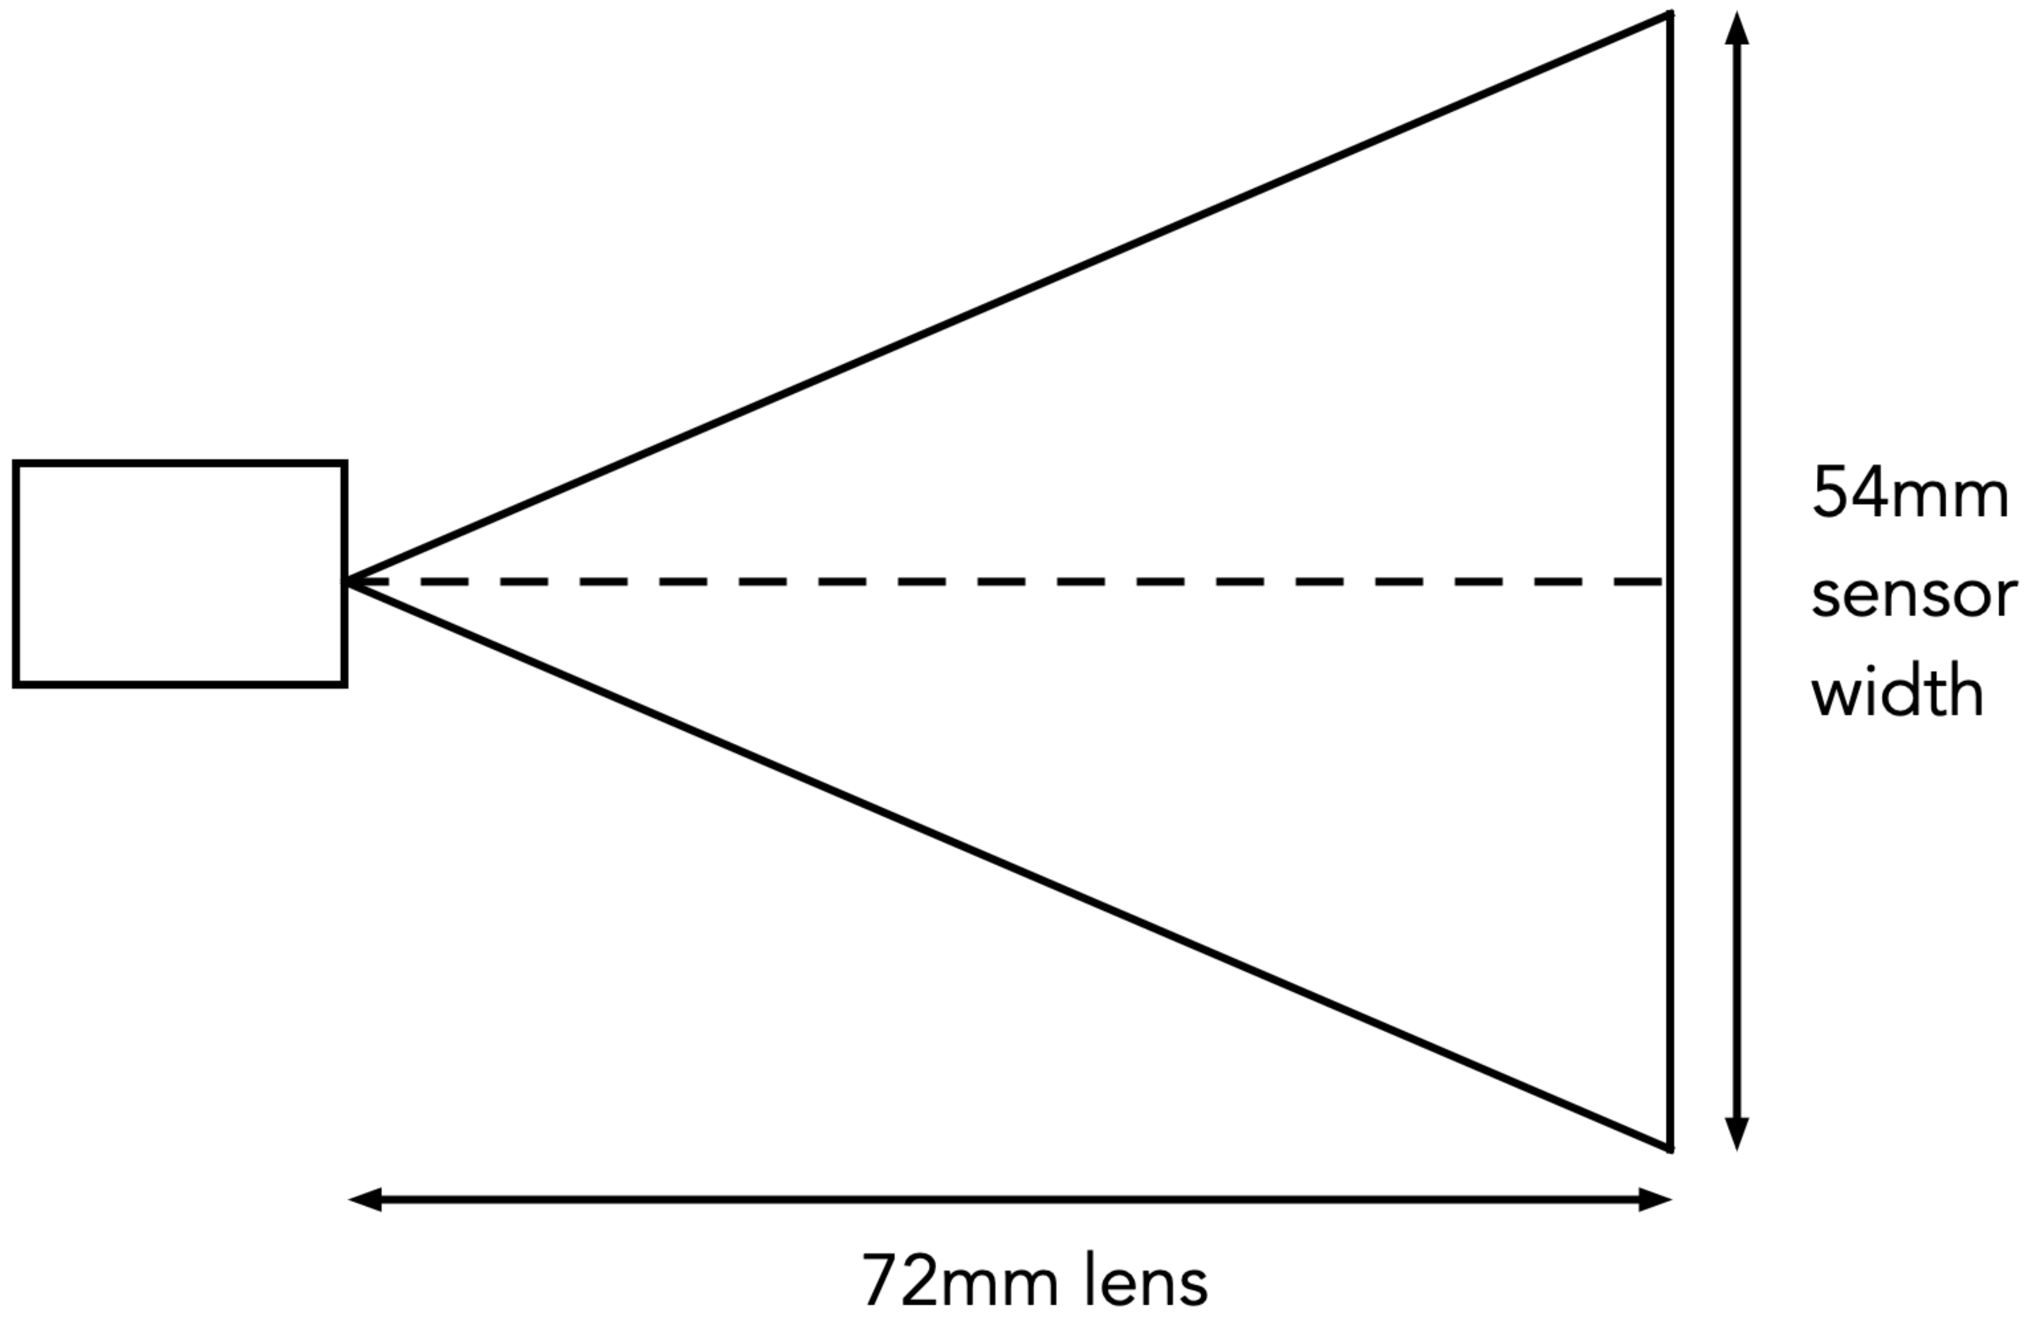

Lighting

- Tested no preferential light direction with a ring light but while it gives a truer colour representation the image appears flat.
- Raking area light tilted at 45 degrees.
  Chosen for aesthetic reasons.
- Raking light reveals surface structure in the painting.
- Lighting is from “the top” to match expectations of sun or ceiling light.
- Professional LED light source with high CRI (colour rendering index) to sunlight spectrum, temperature of 5600 Kelvin.
- Cross polarisation employed to remove specular reflections.

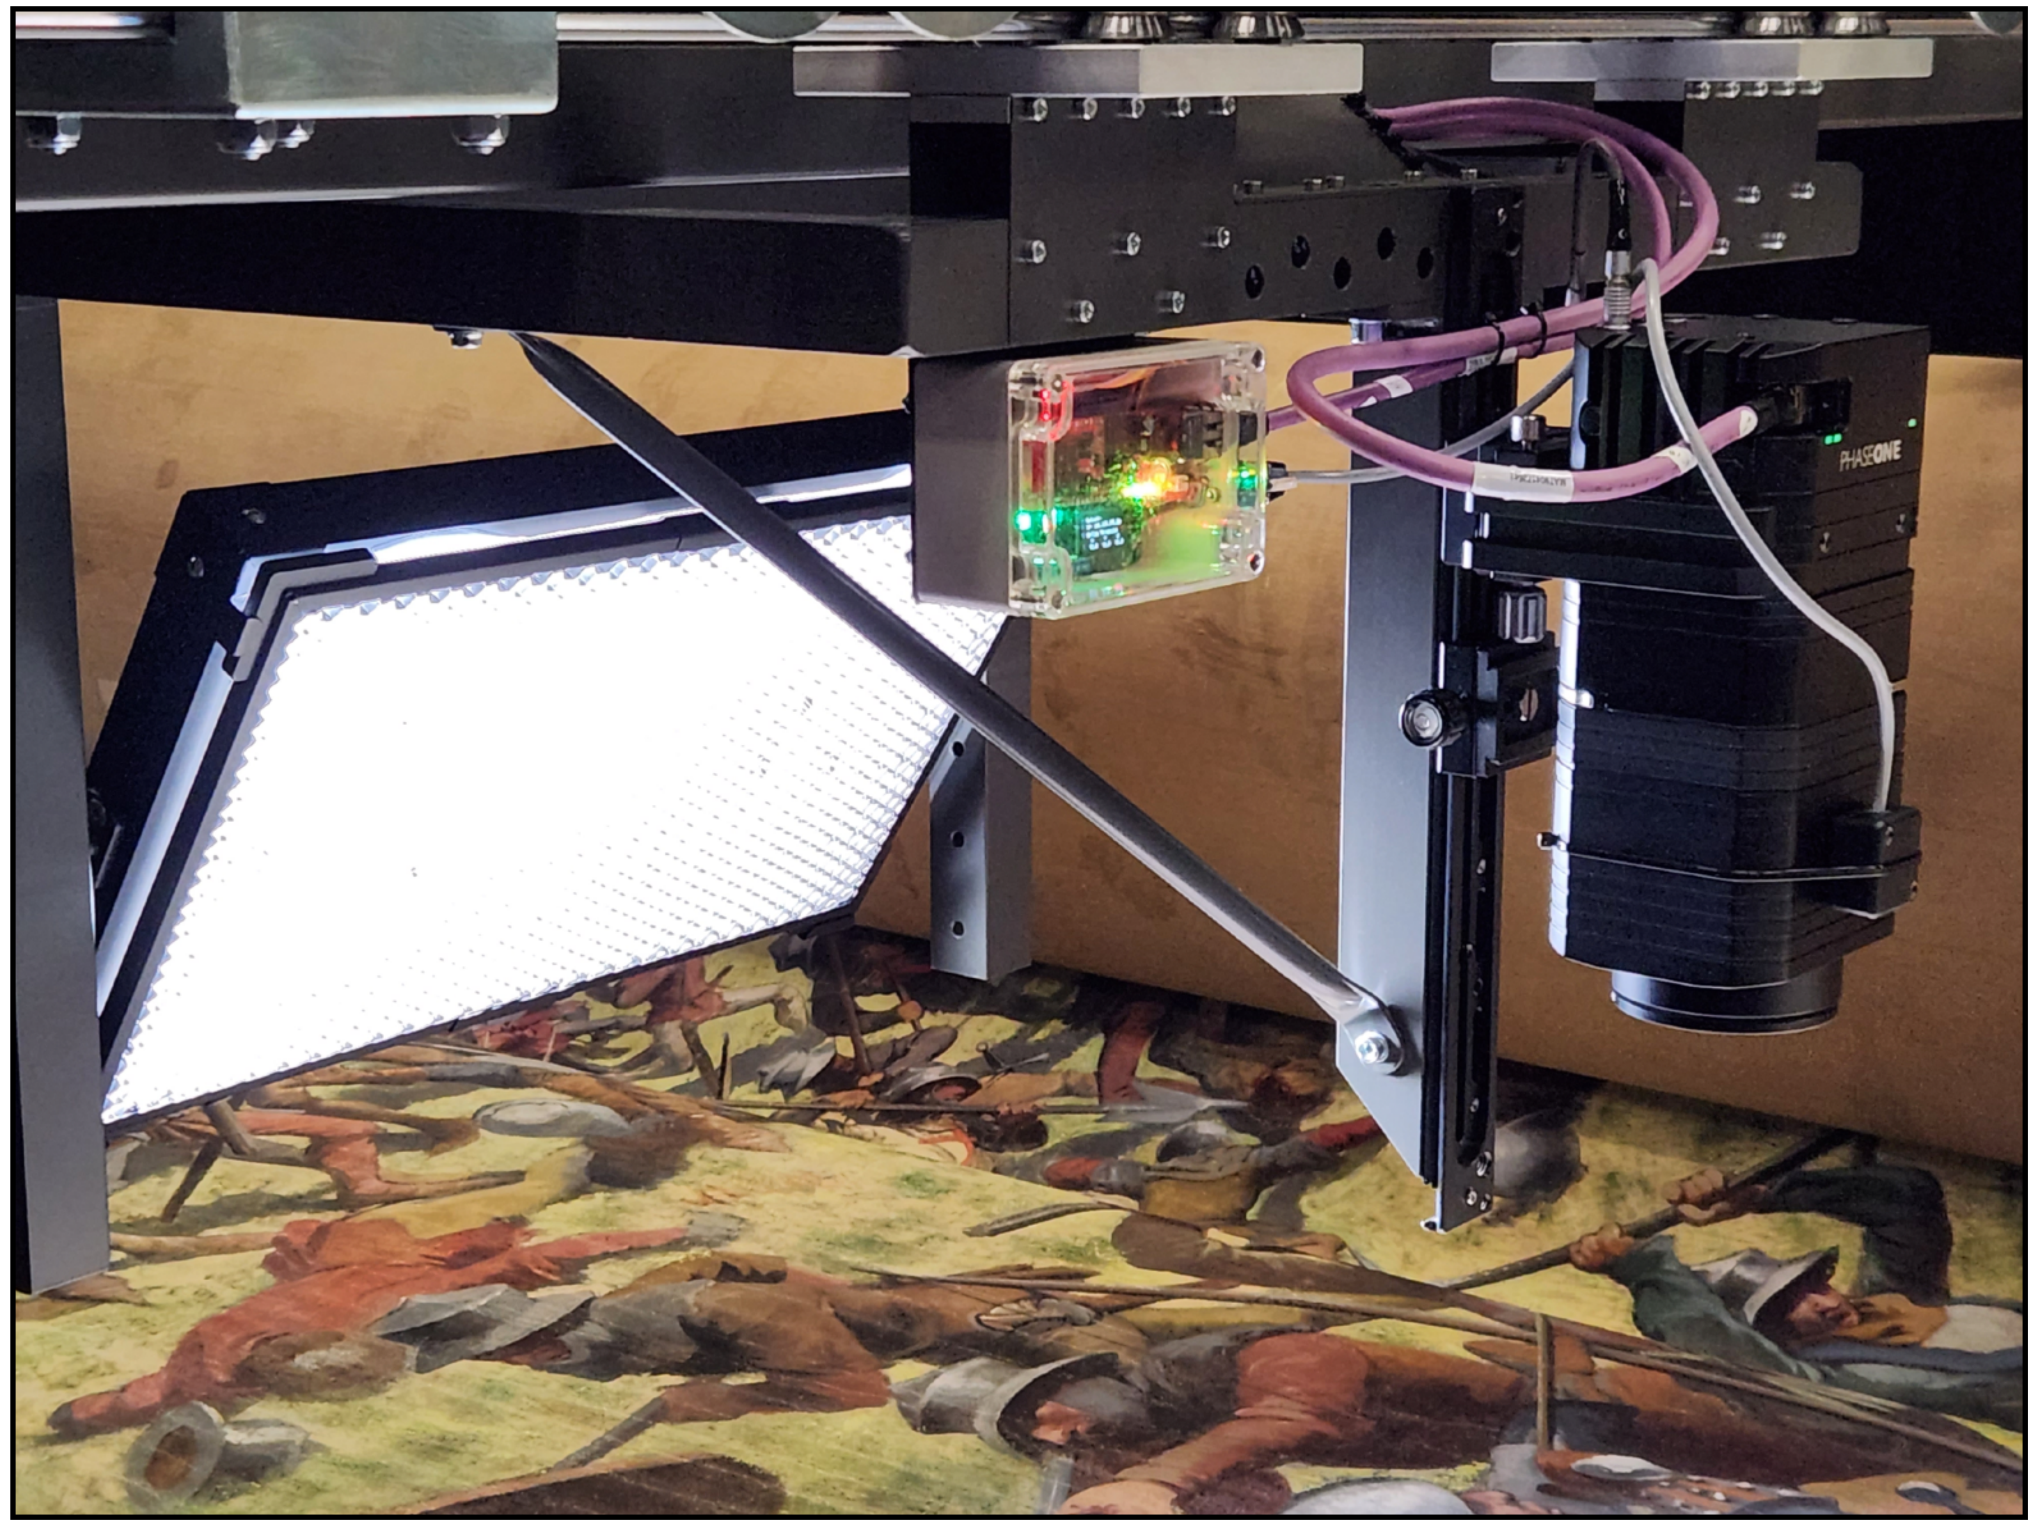

Raking area light

Initial lighting experiments

Stitching tests, July 2022

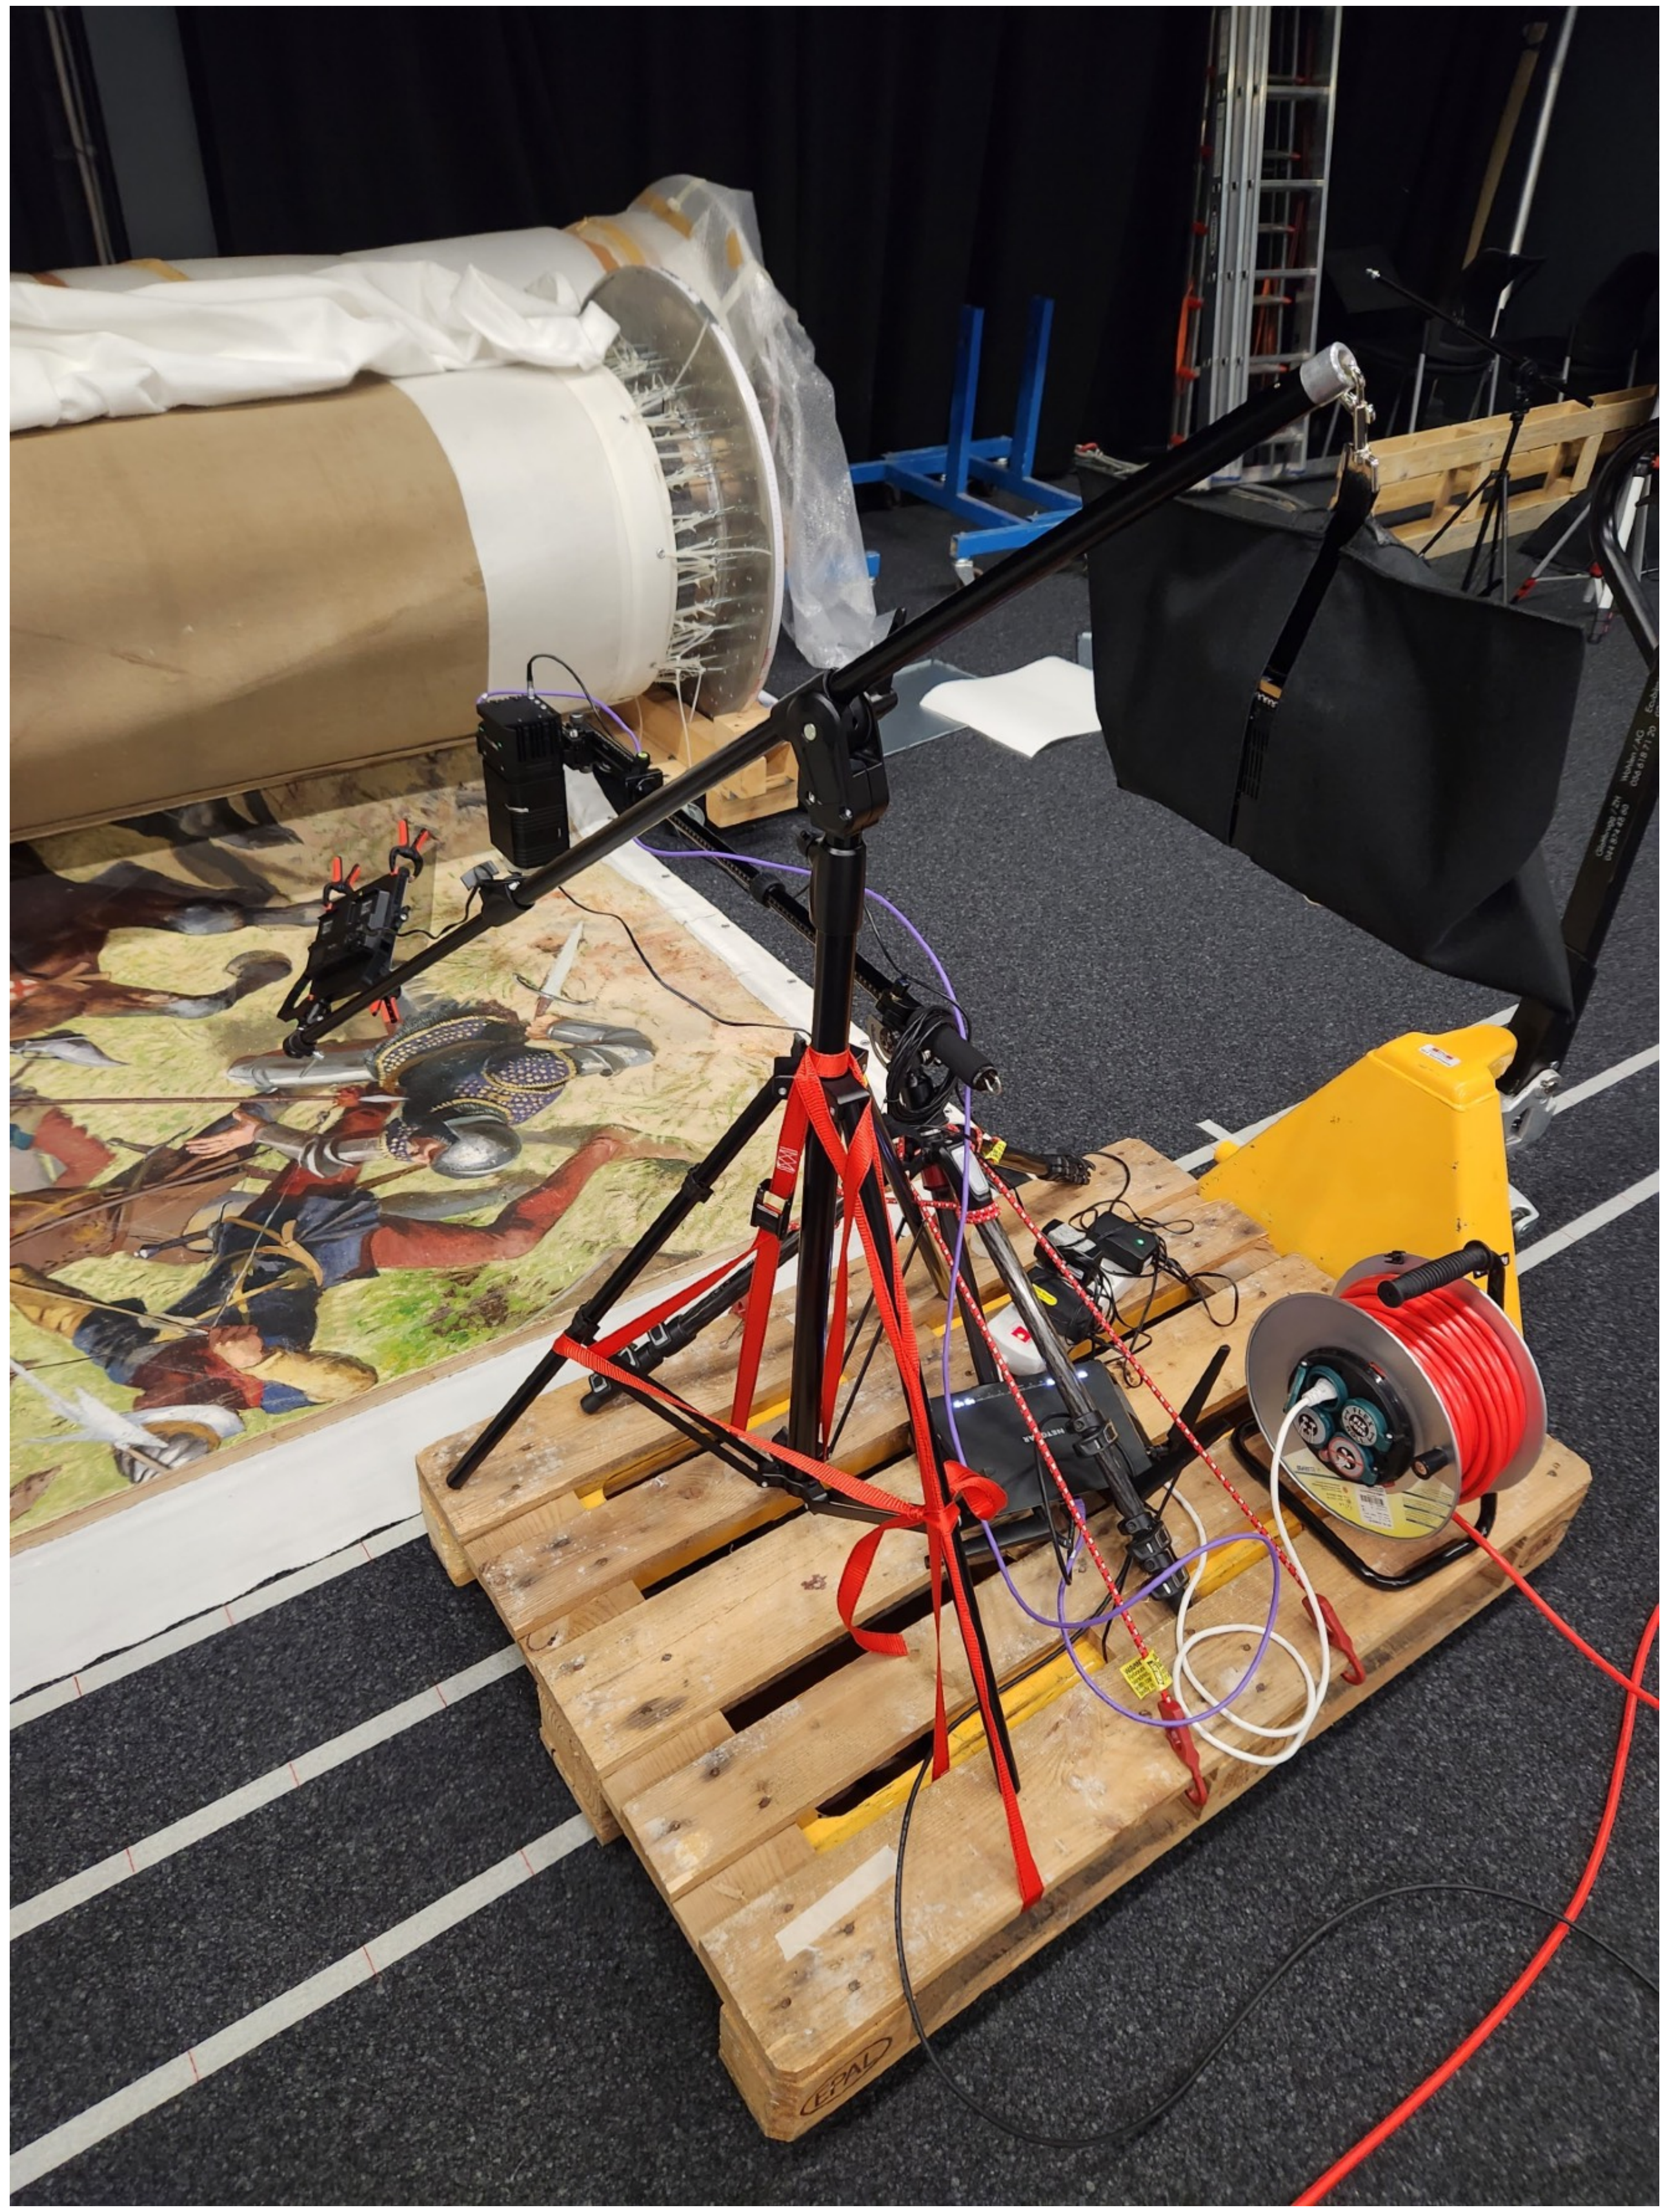

Lighting tests, November 2022

Cross polarisation

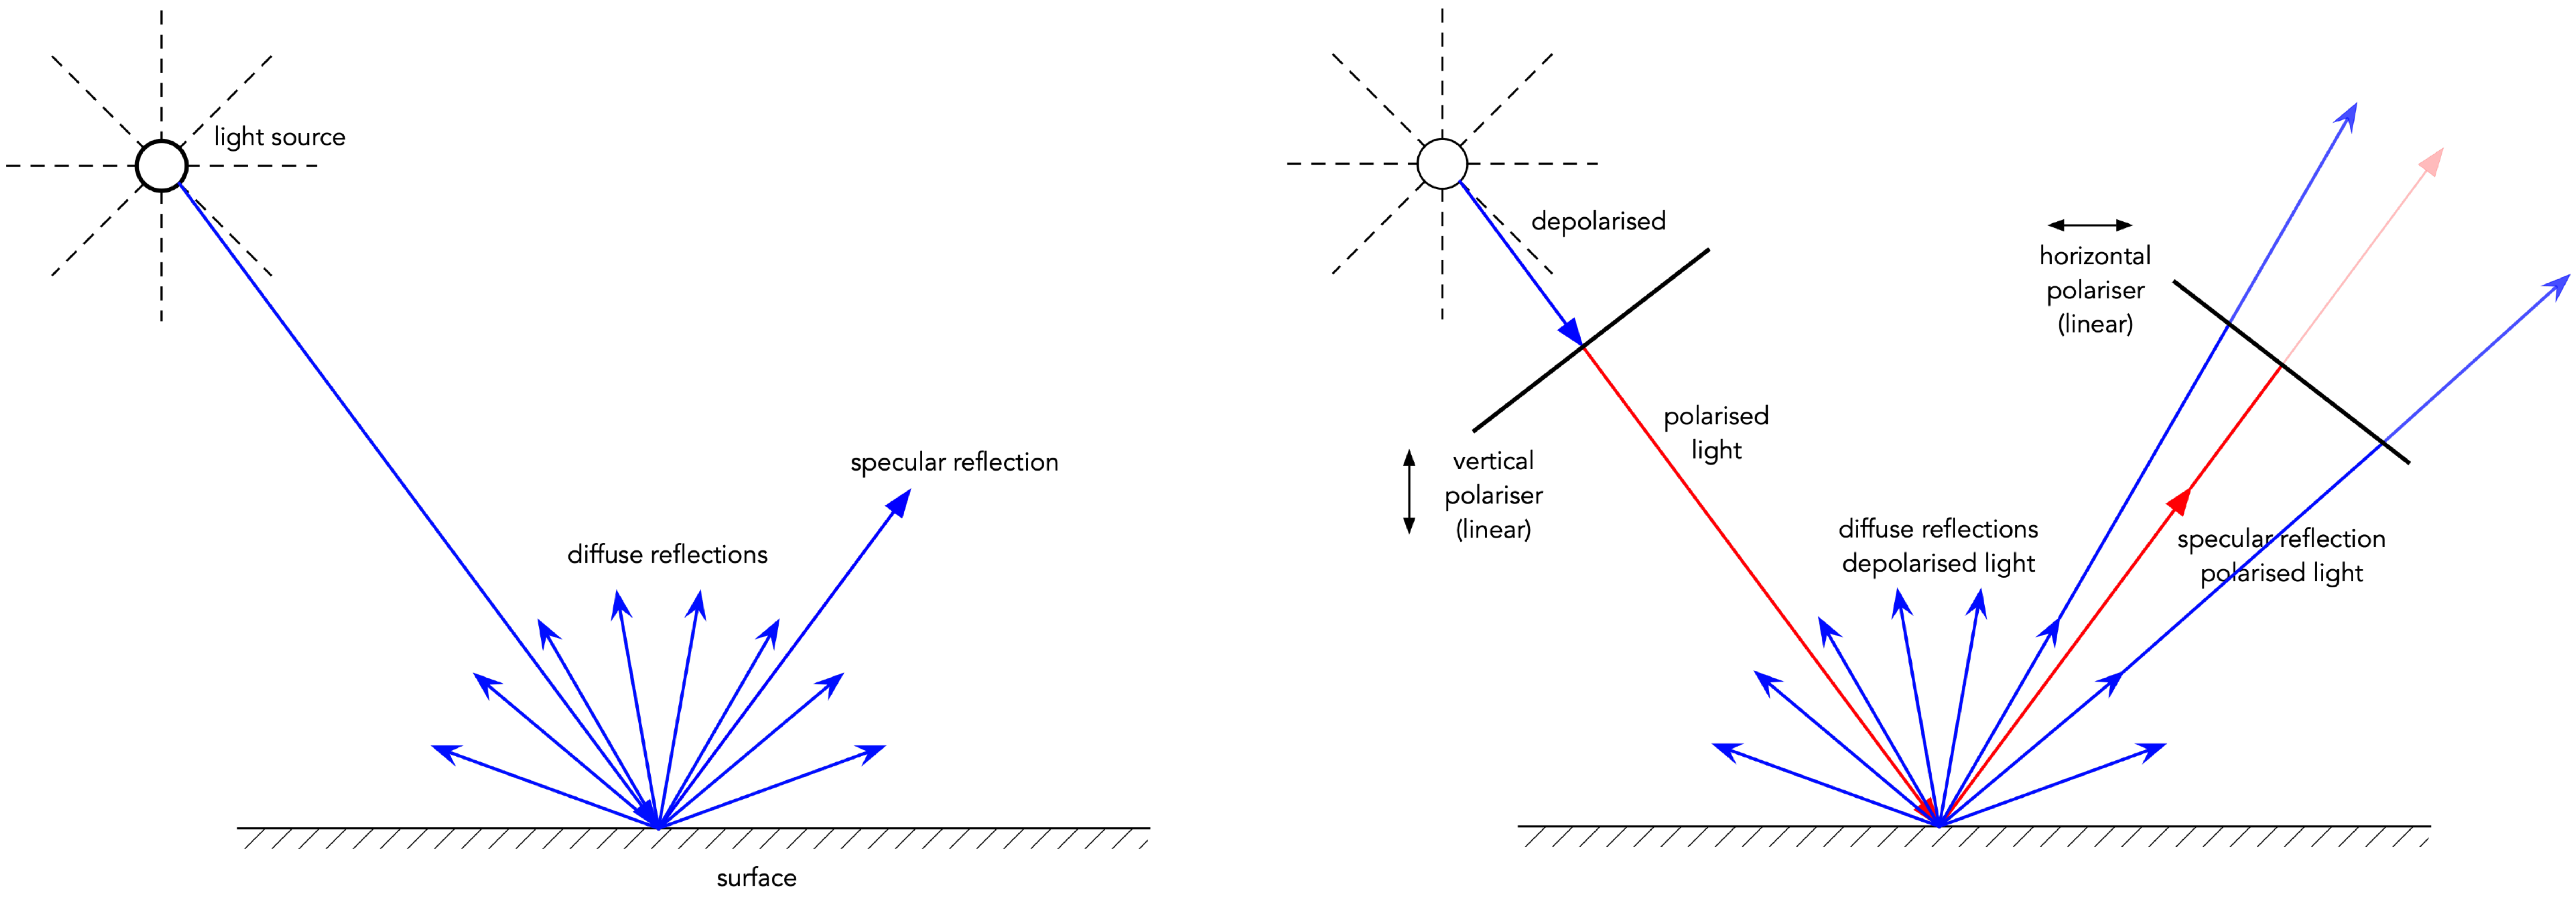

Cross polarisation comparisons

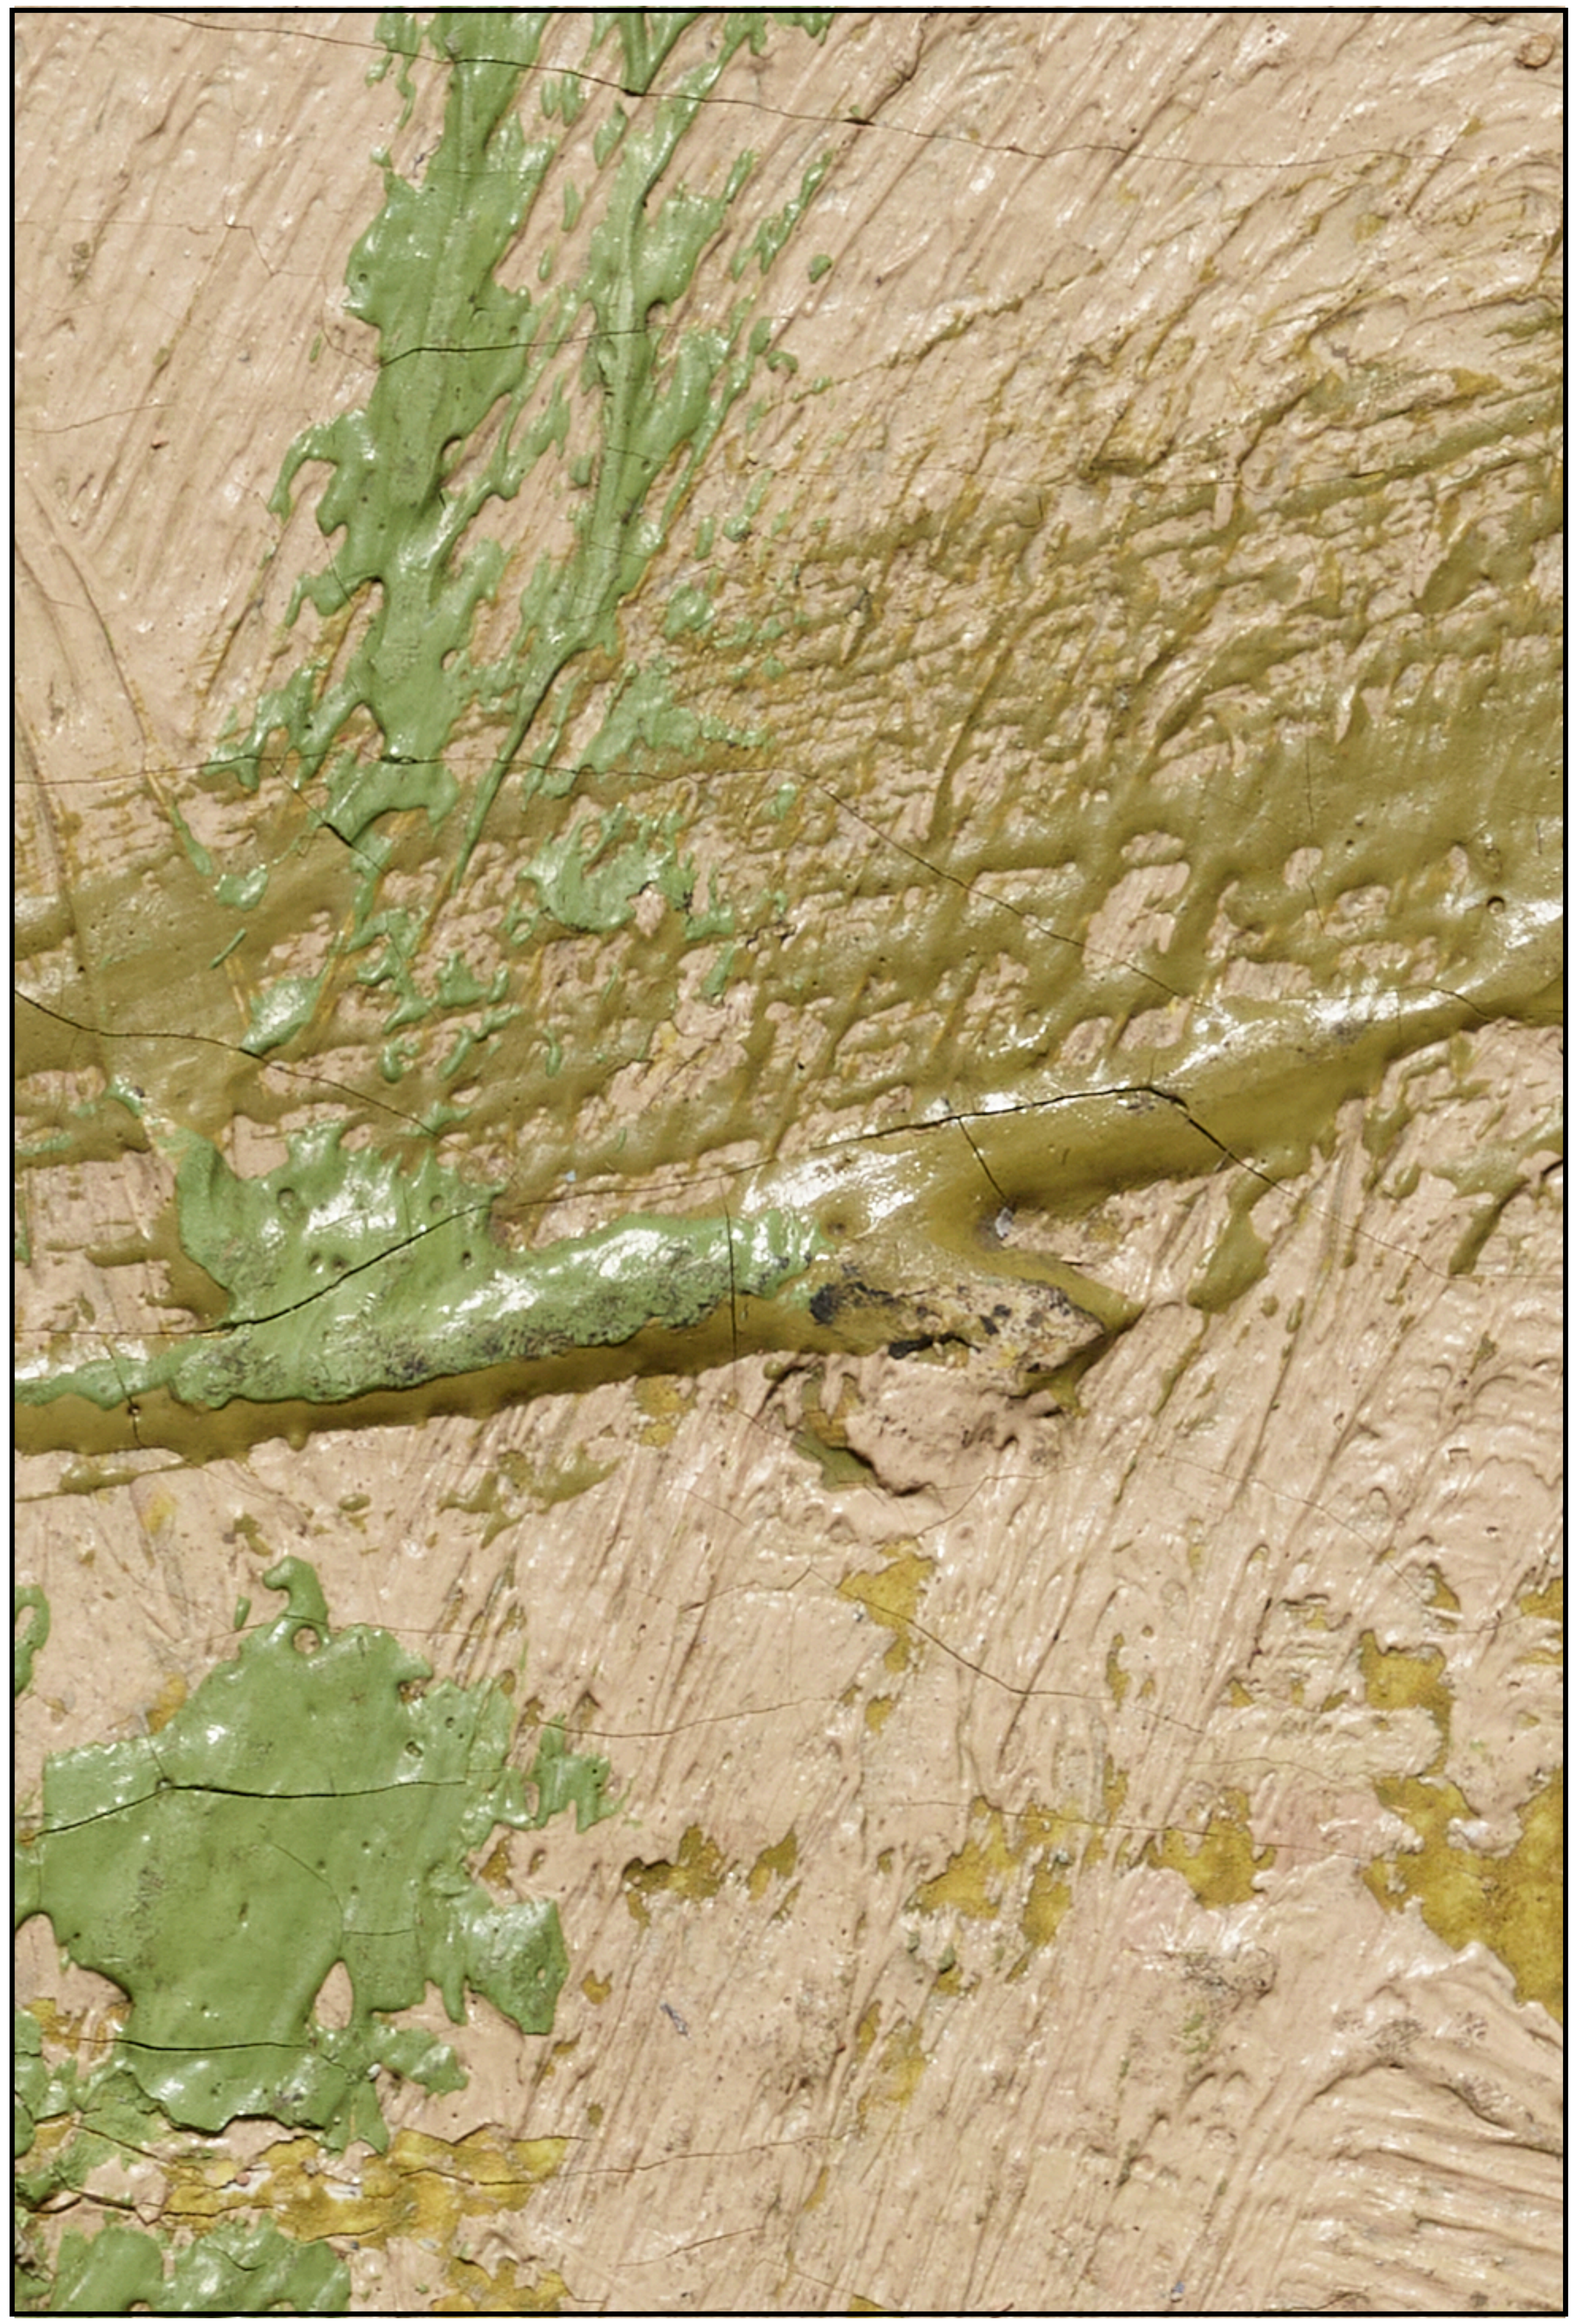

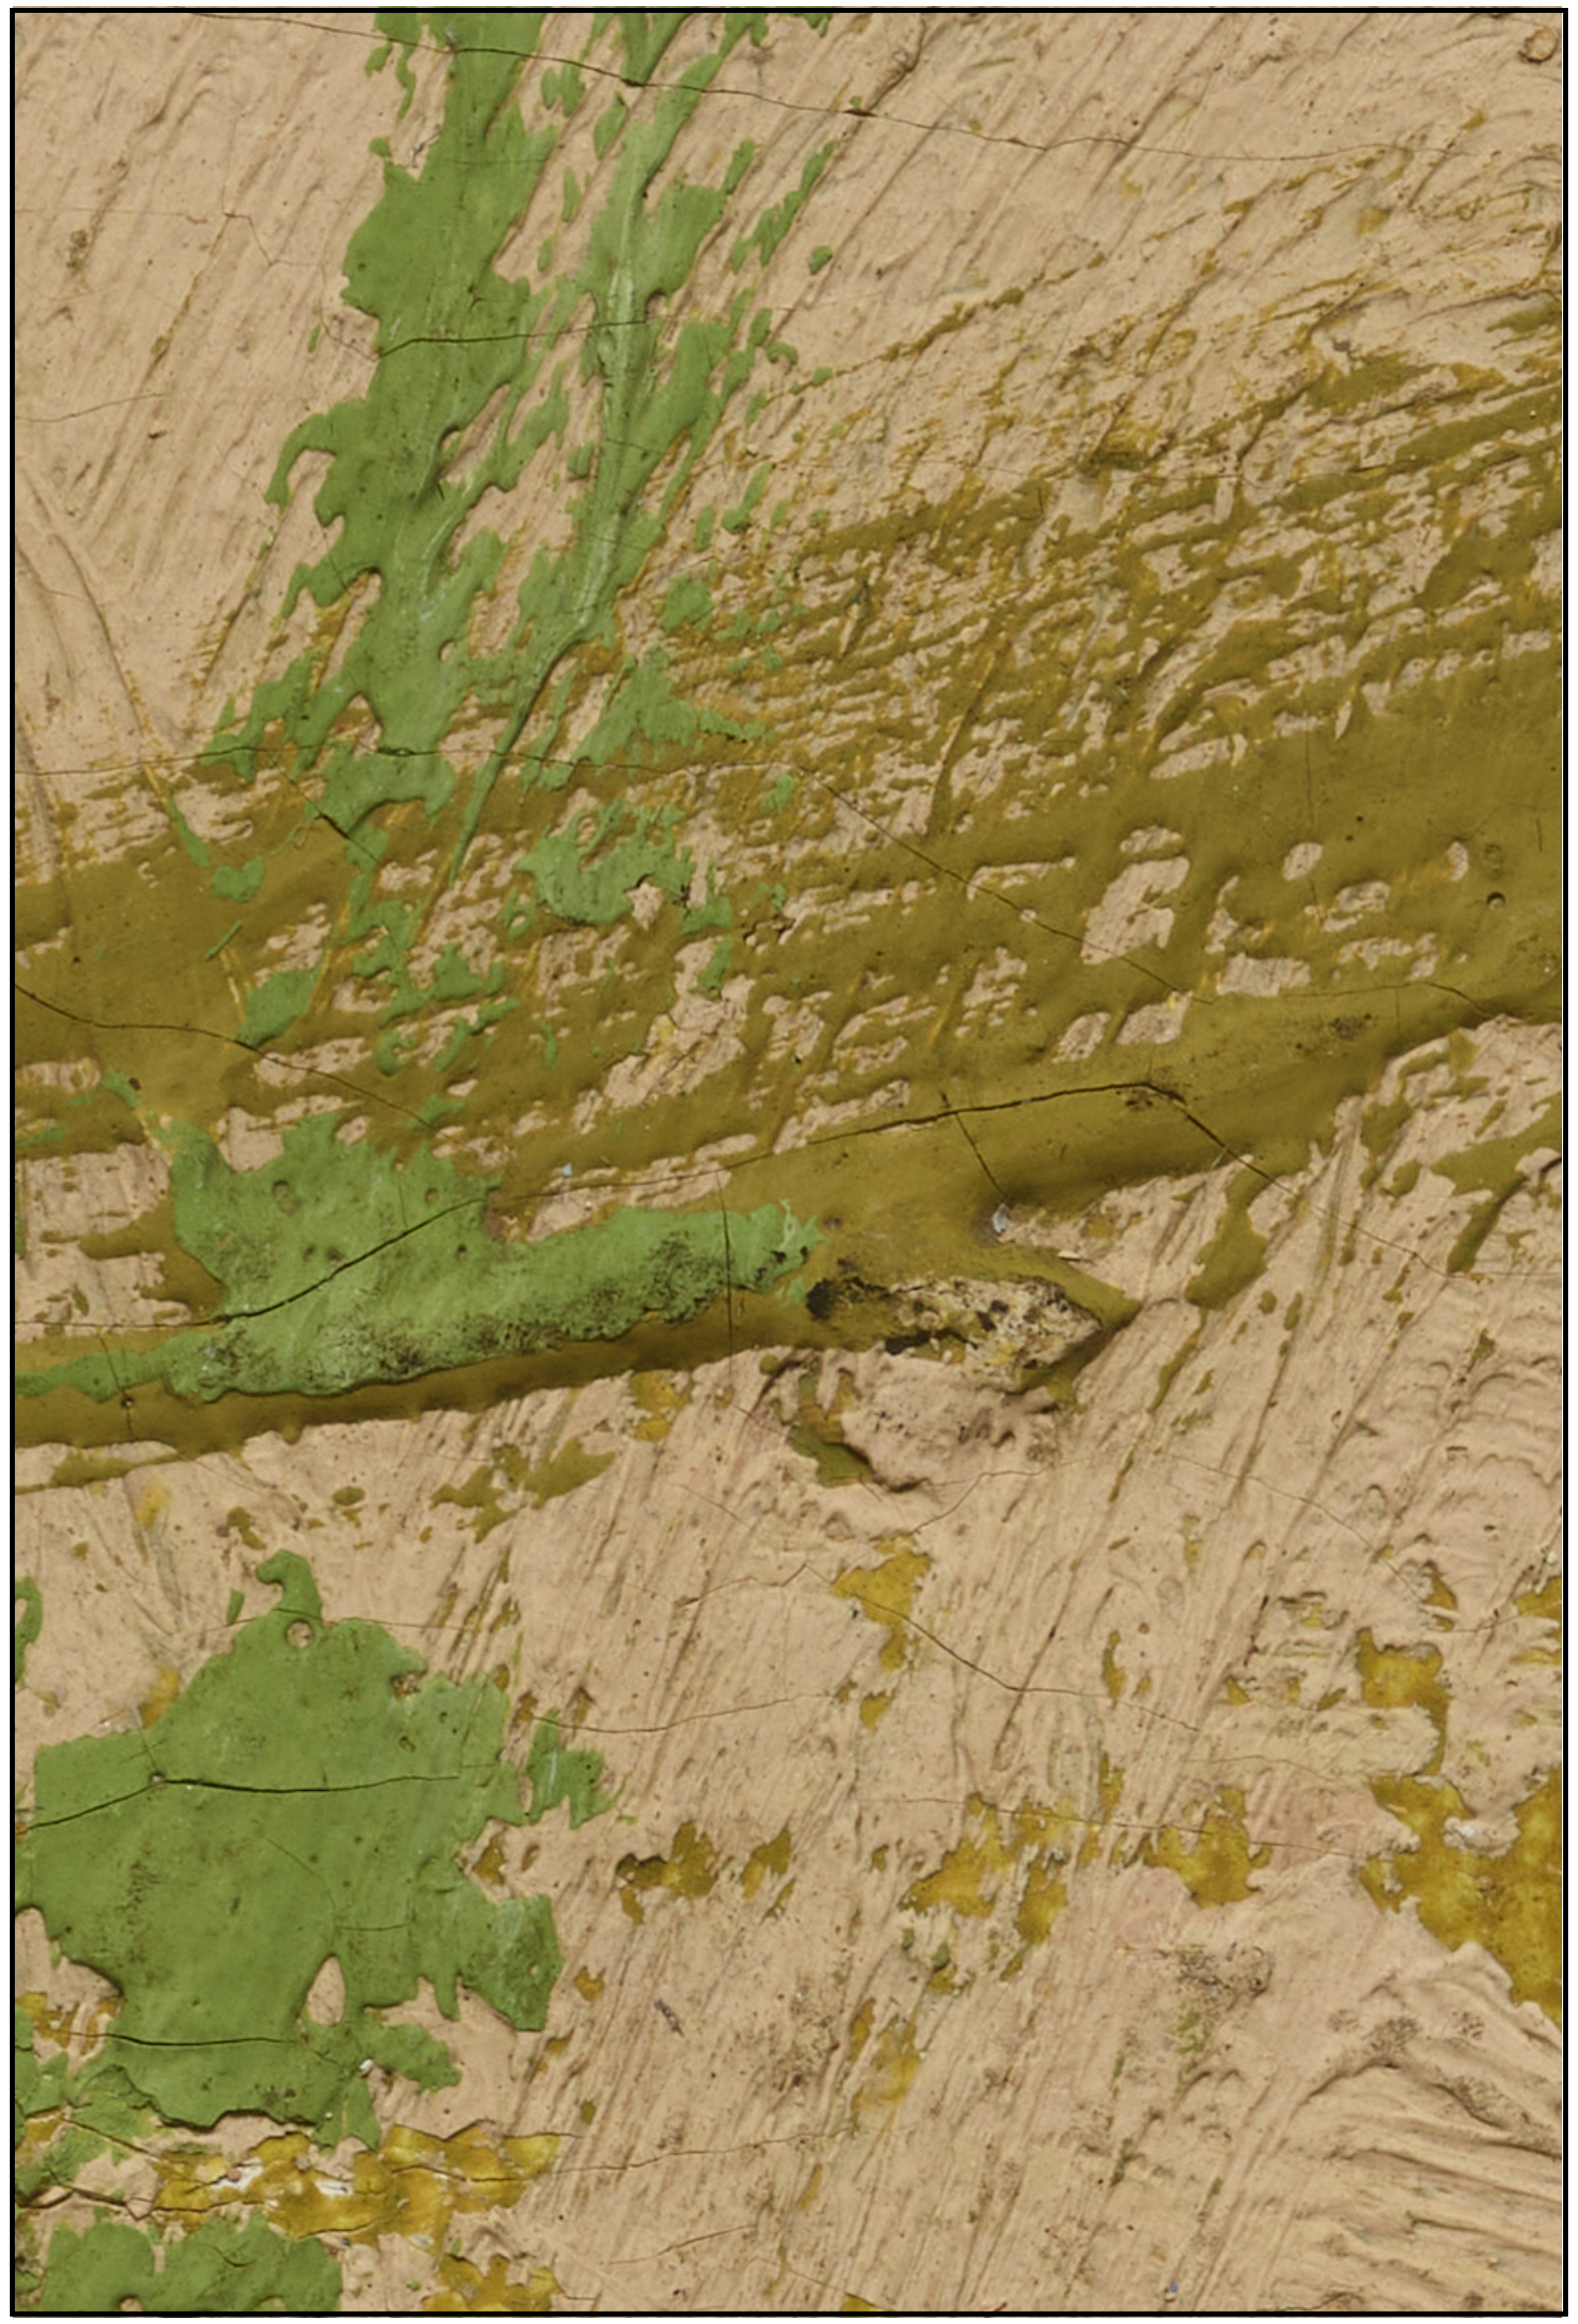

Cross polarisation comparisons

Cross polarisation: metals

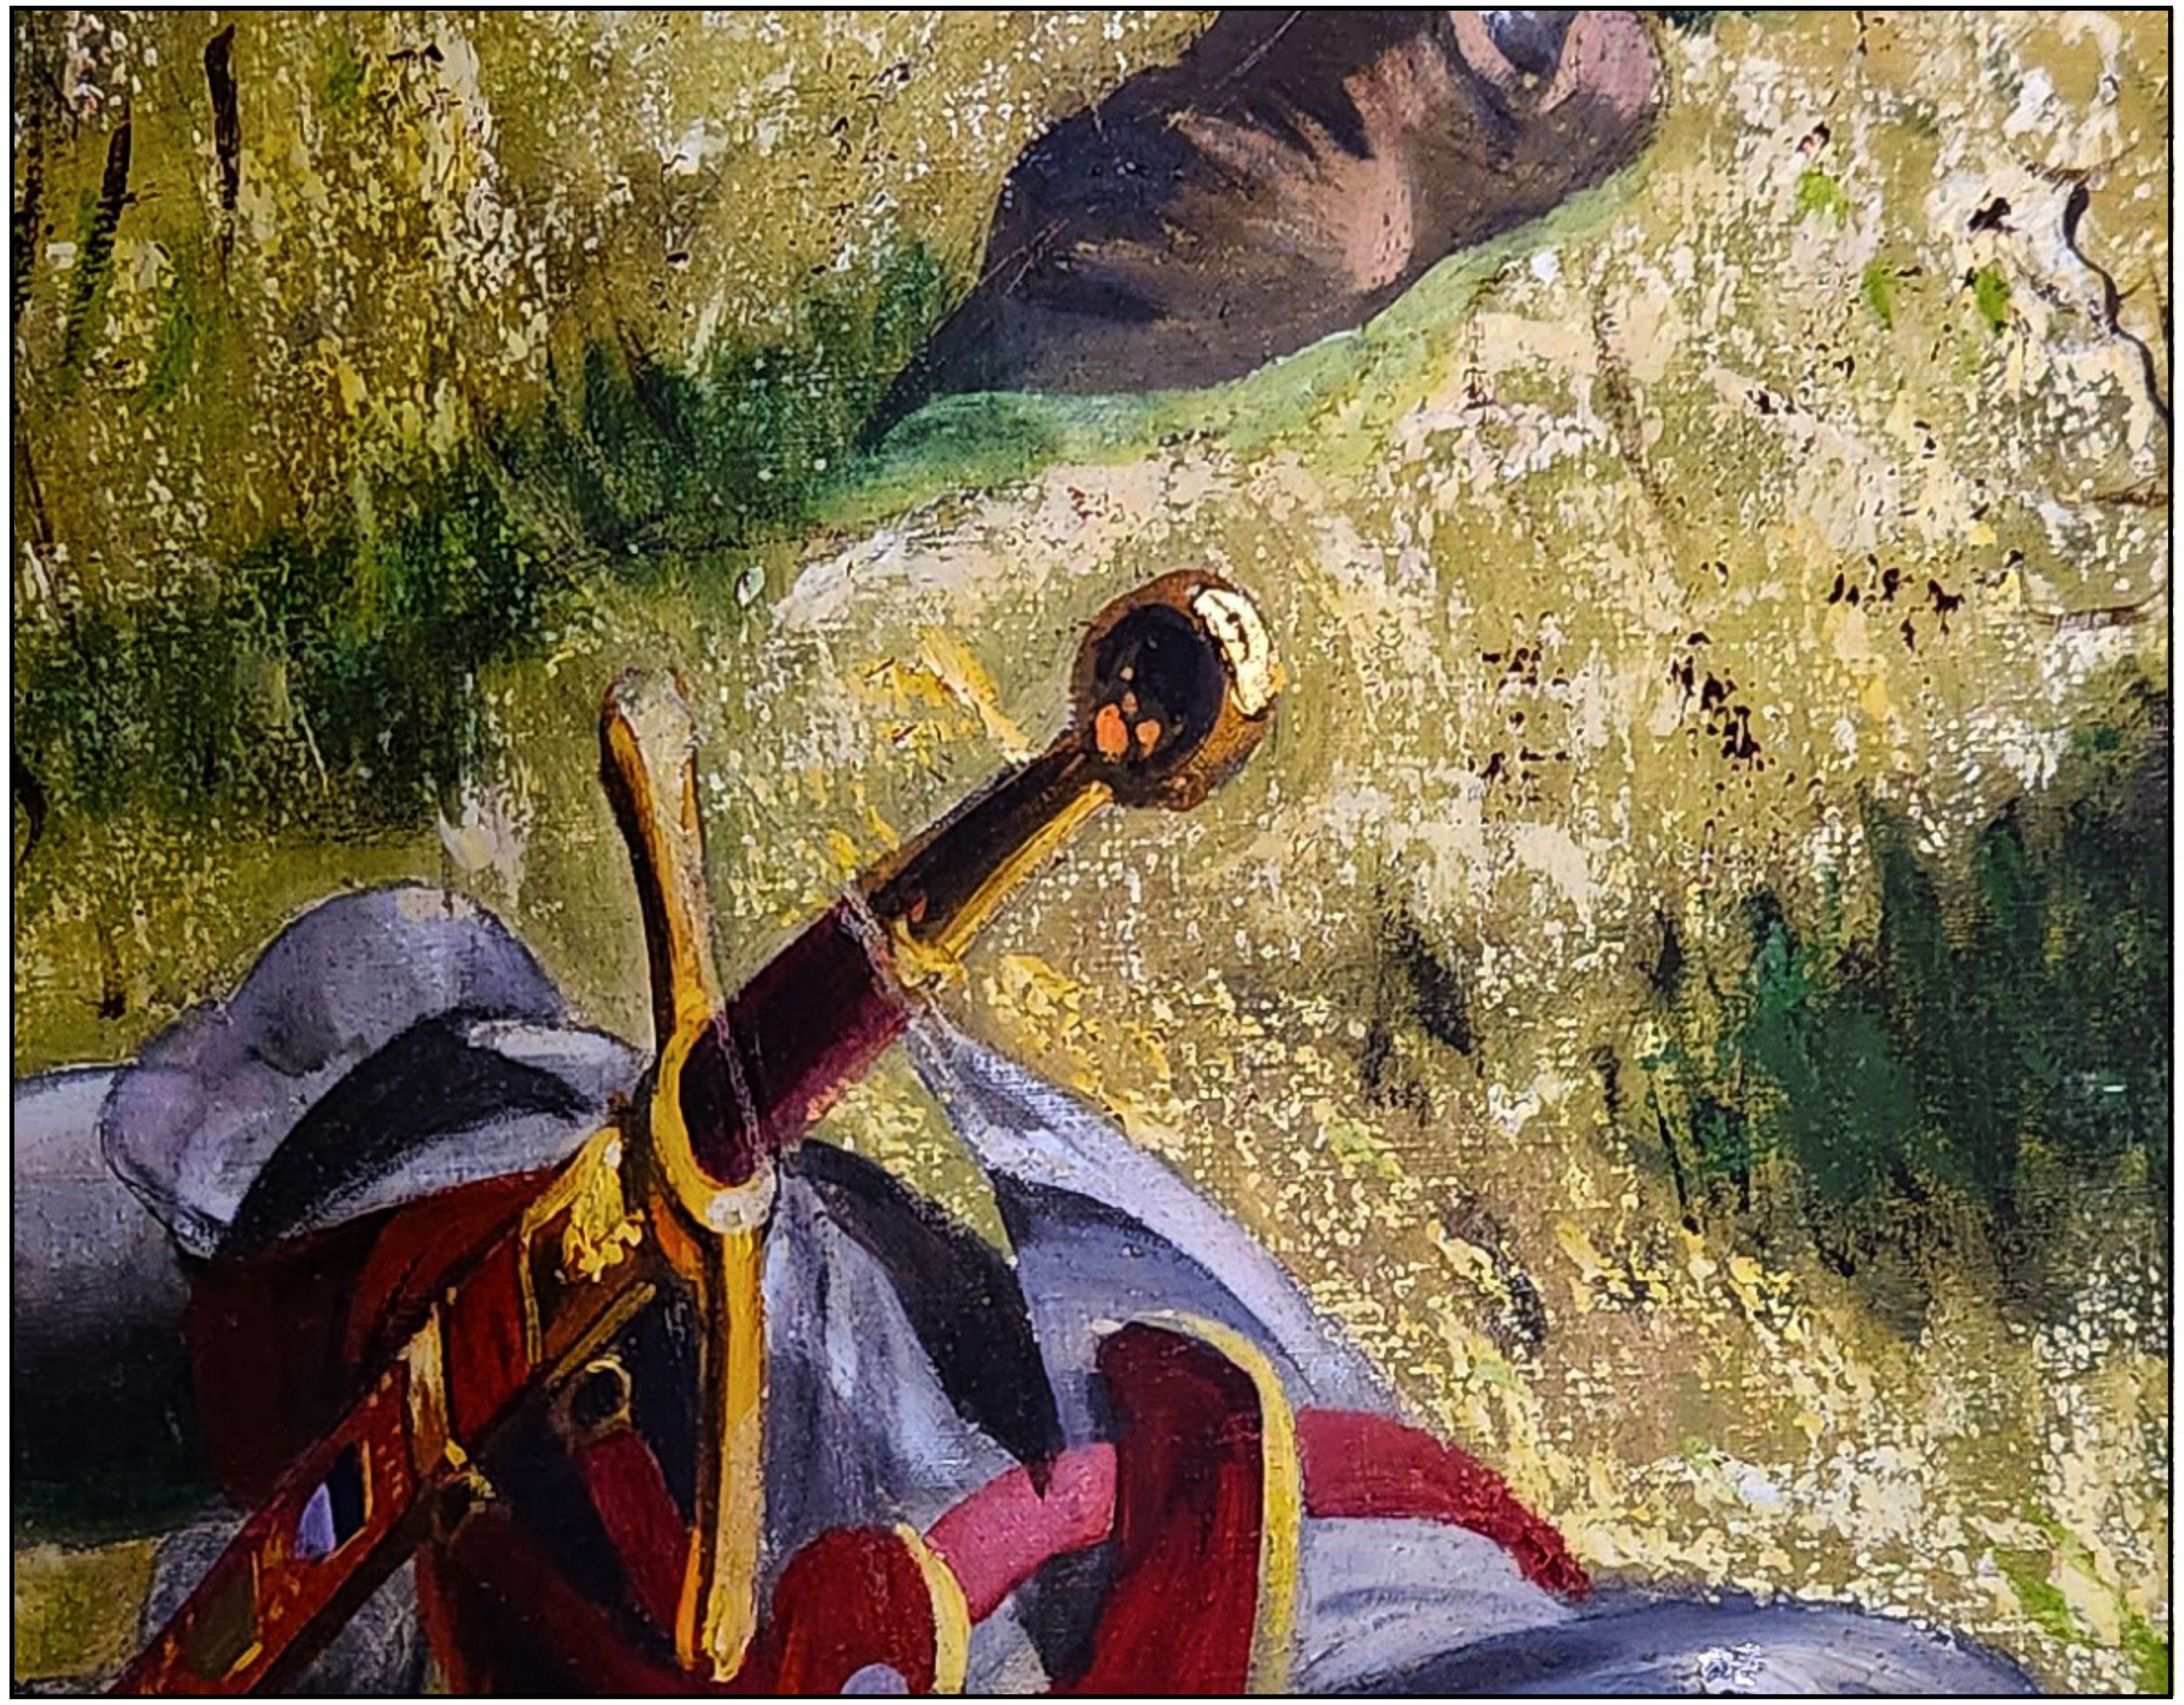

Cross polarisation: metals

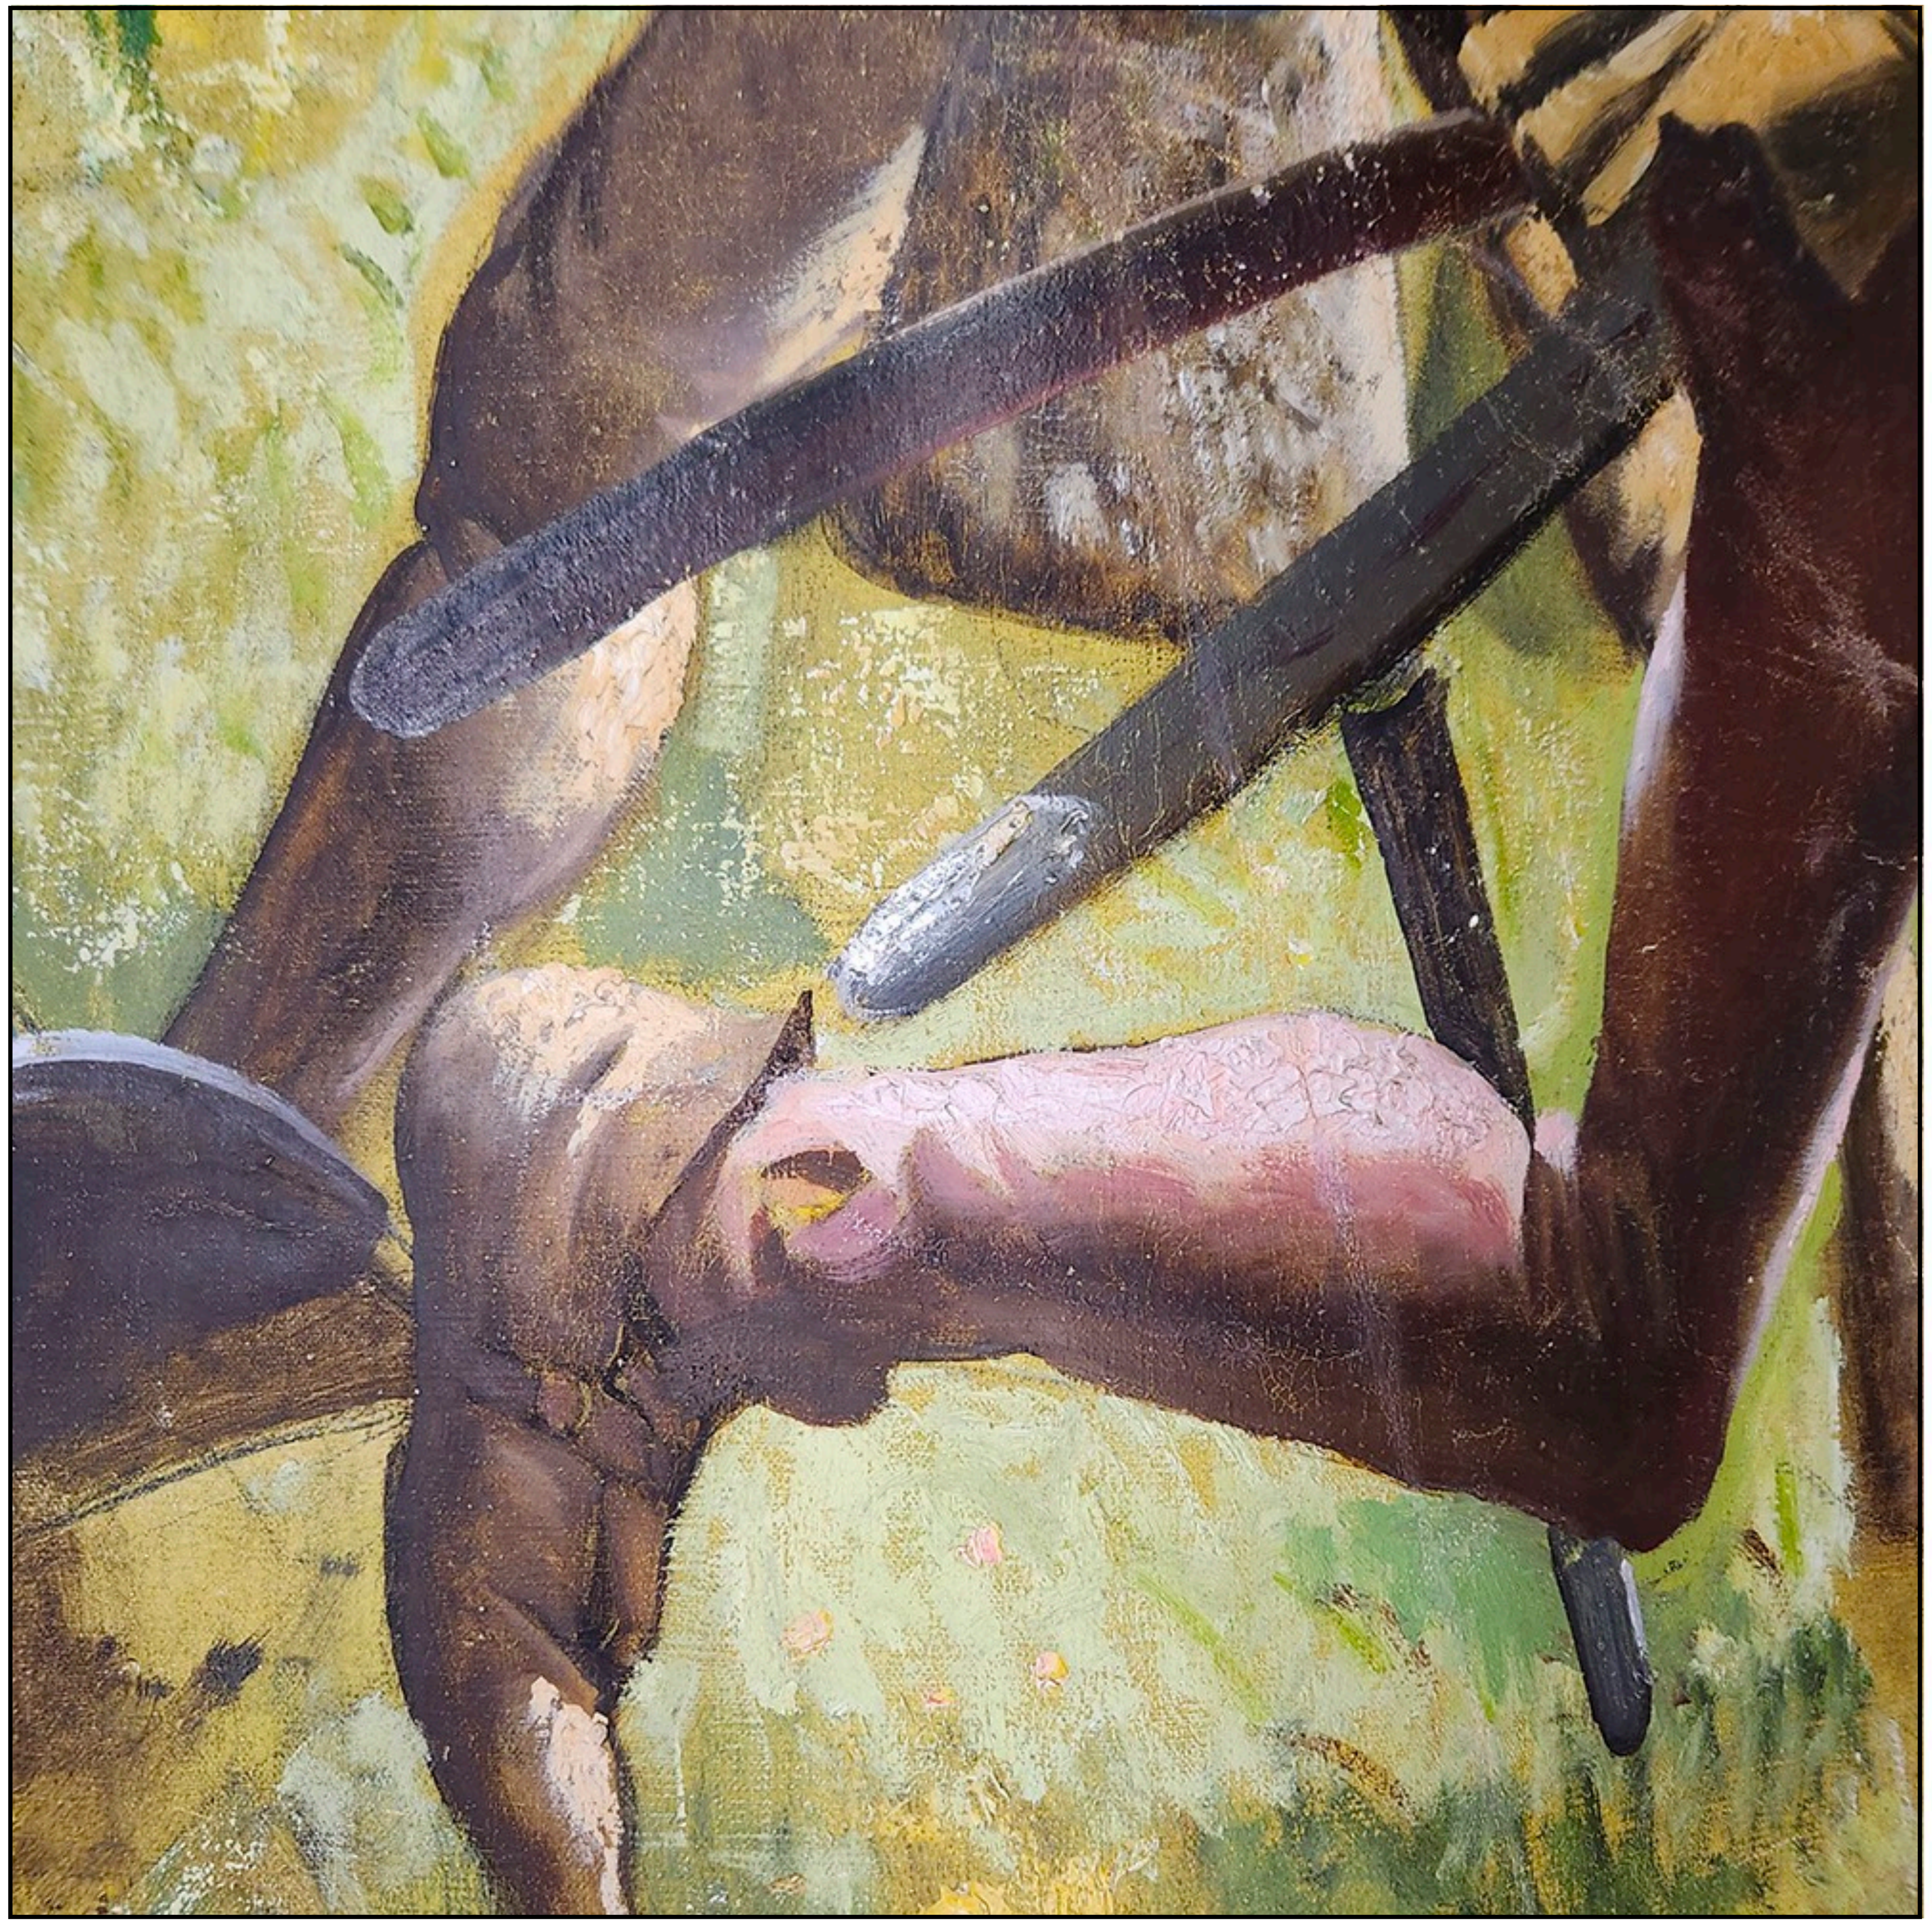

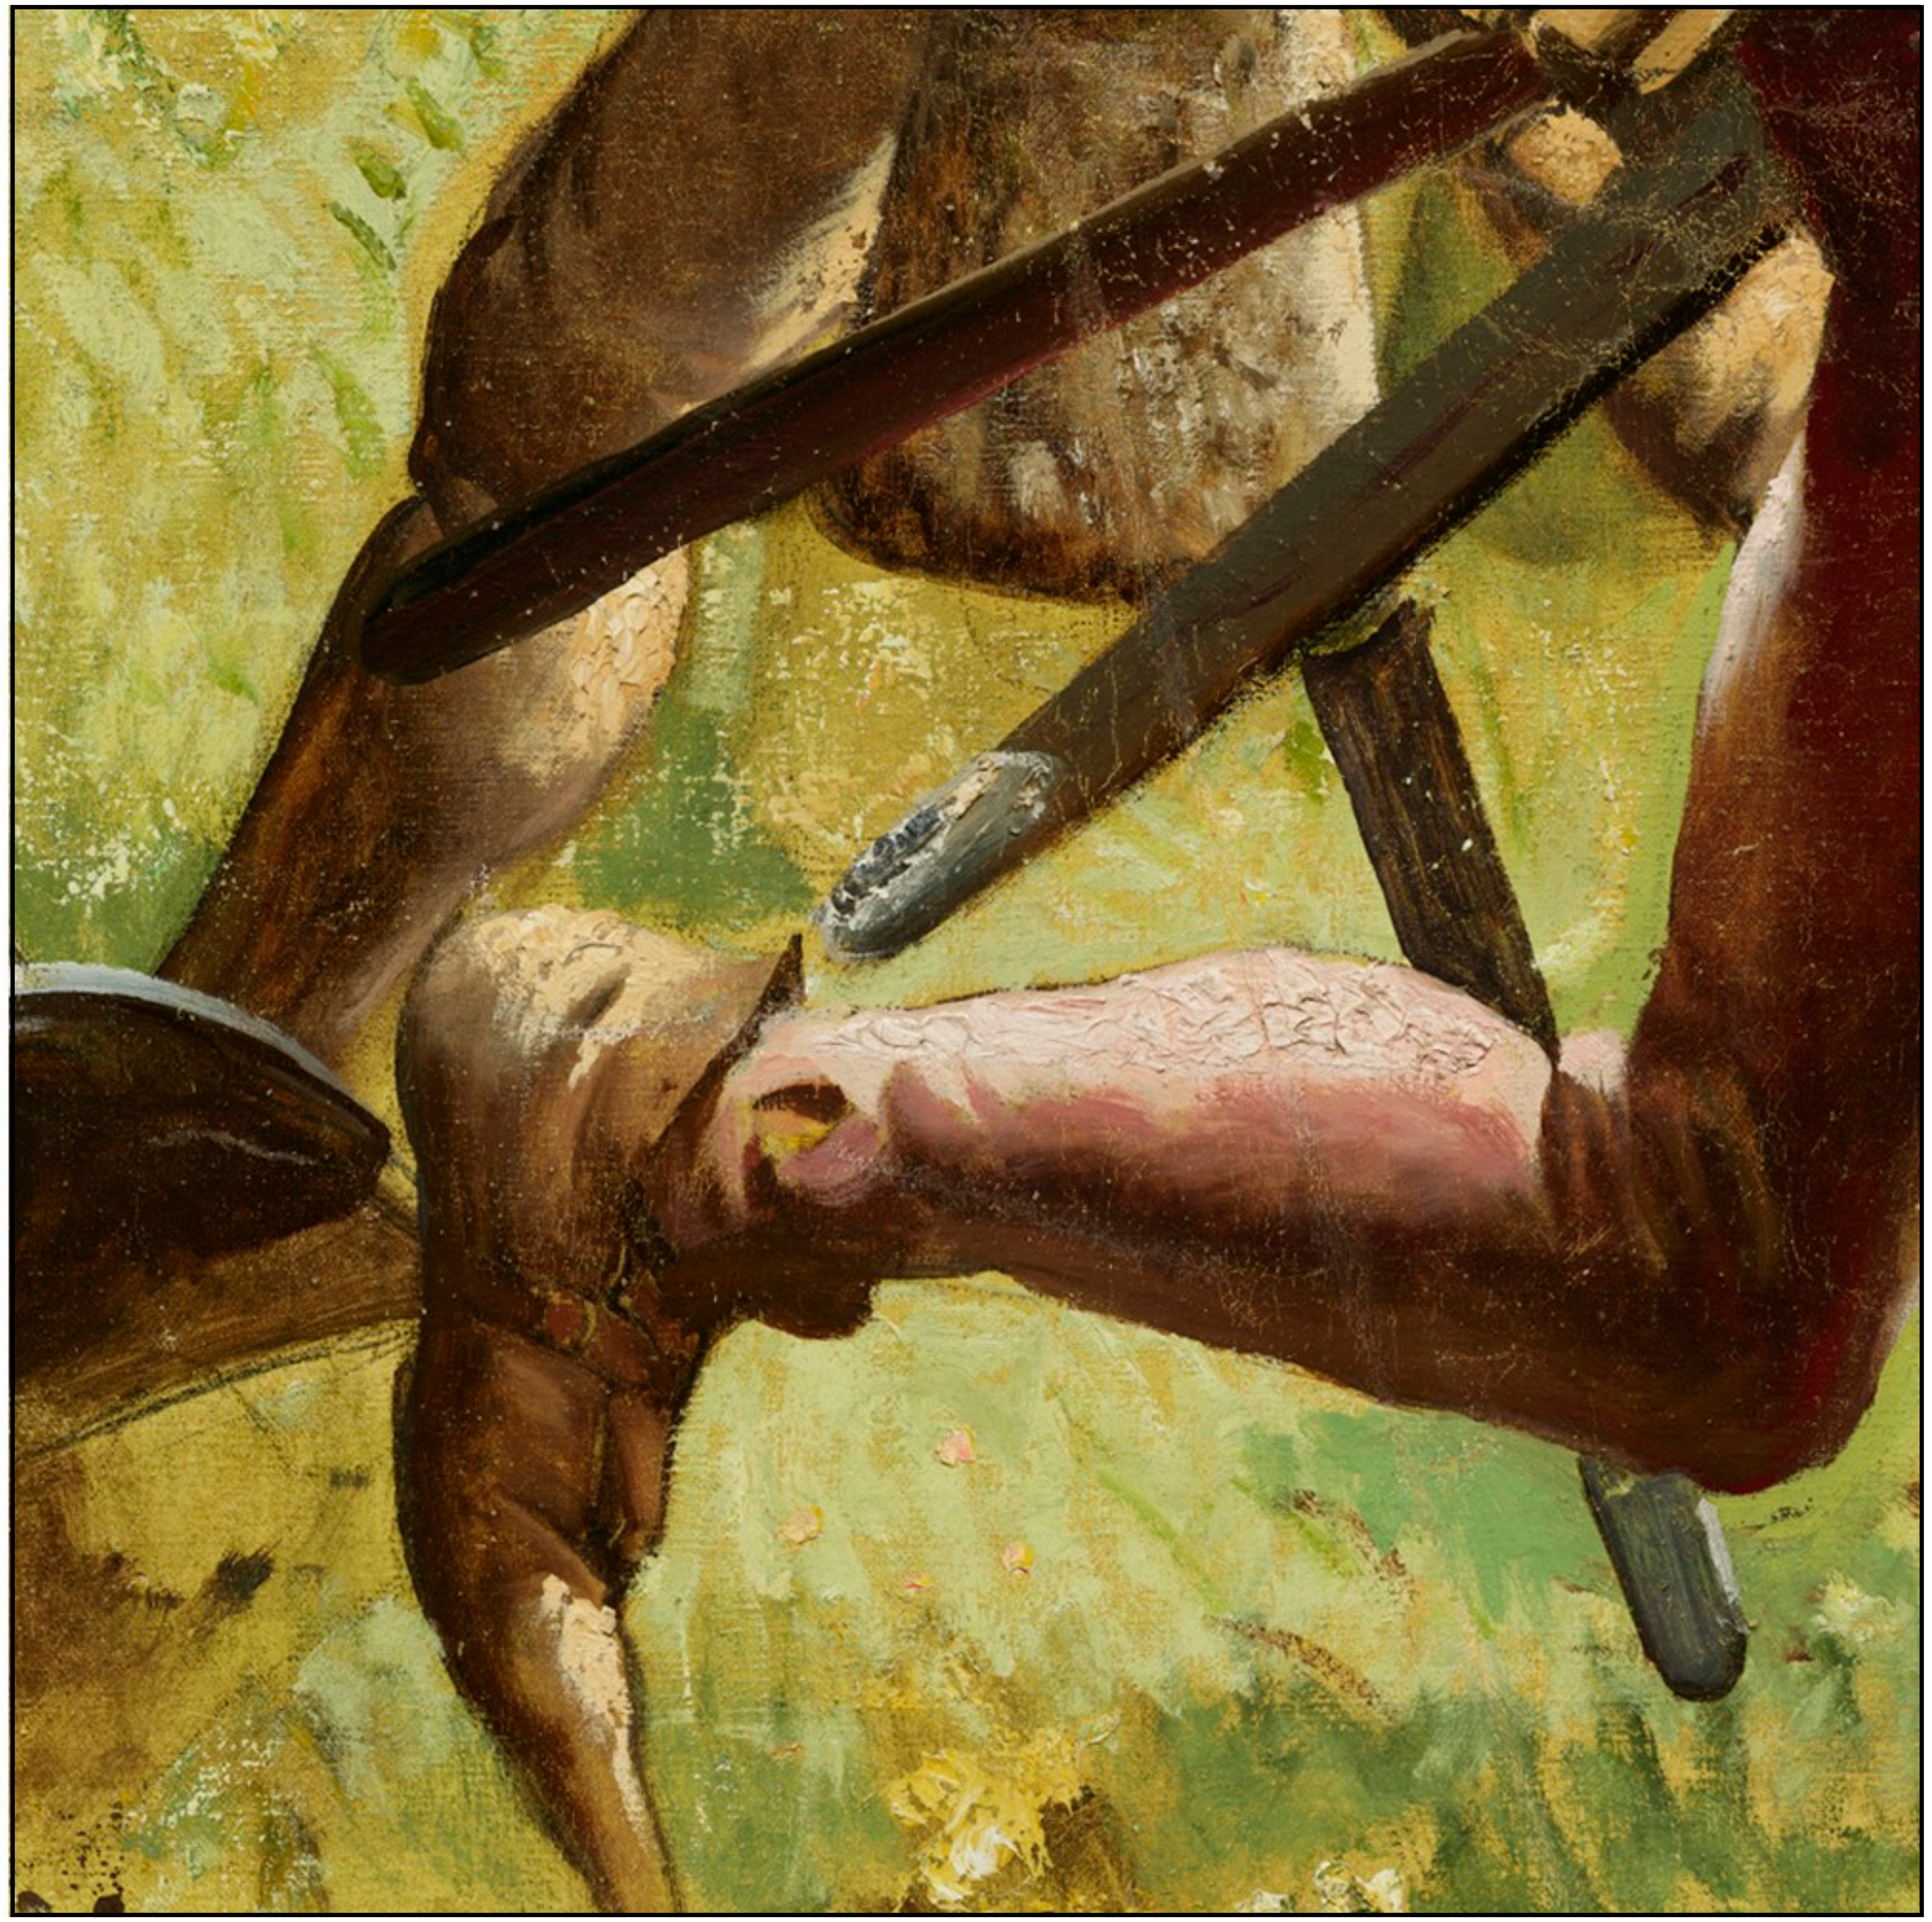

Hyperboloid

- A surface tensioned between two rings forms a hyperboloid, a shape in 3D space.
- It cannot be laid flat physically.
- The painting cannot be perfectly represented as a 2D rectangle.
- The inner circumference is less than the upper or lower circumference.
- The camera lens optics only has about +/- 1cm depth of focus, so require a flat surface.
- Focus stacking was deemed impractical at this scale both in terms of time and resources.

Influenced almost every aspect of the project.

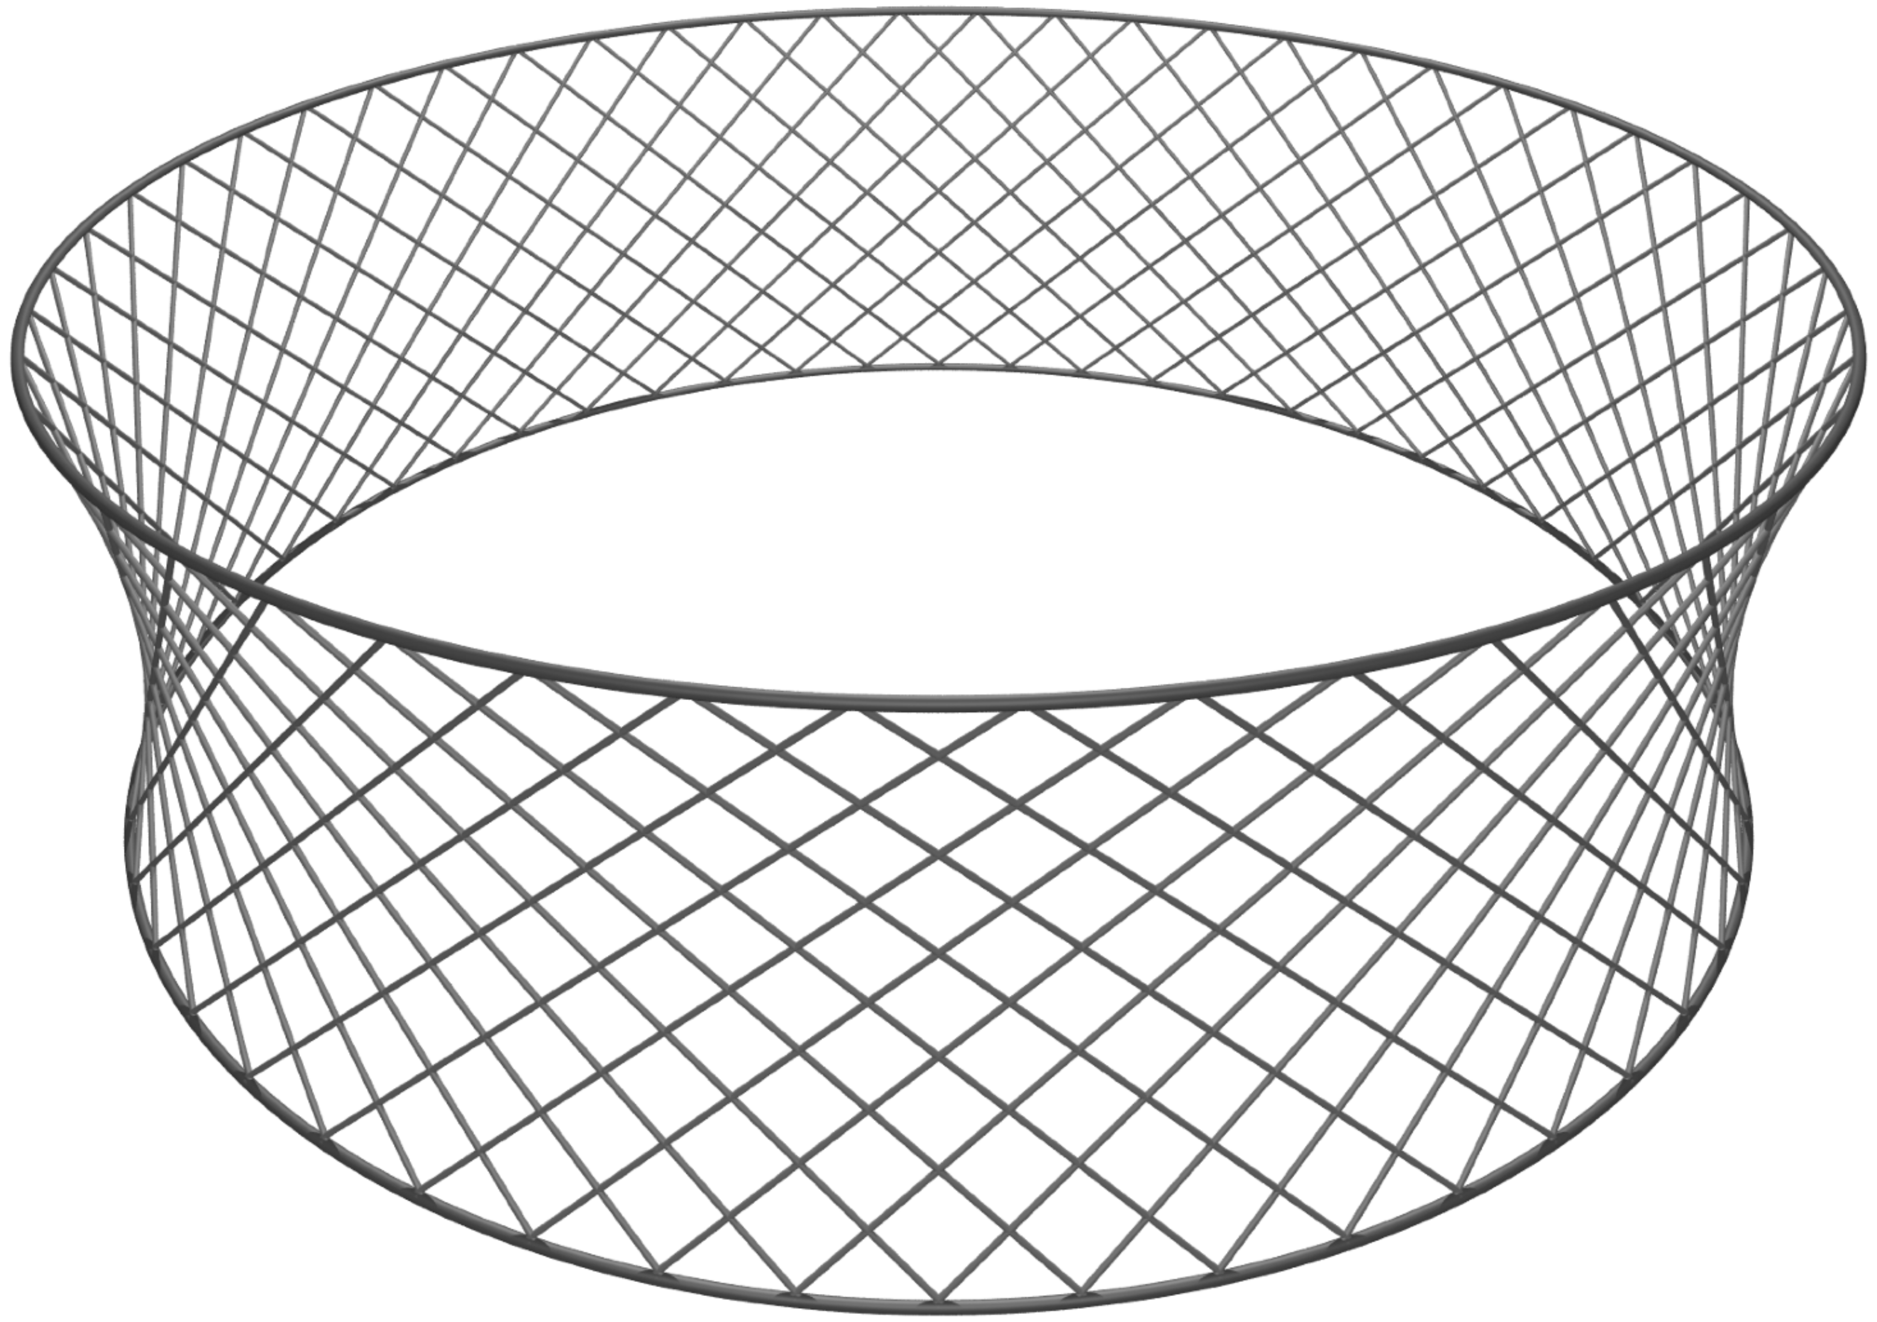

Cross section of 2002 exhibition building shows hyperboloid nature

Resulting folds if laid flat

Scanning rig, 1m wide strip lays flat under tension of the canvas on the sides

Software

- CaptureOne DH, commercial software from PhaseOne.
  - Scripted using AppleScript to control the camera
  - Perform flat field and colour correction
  - Export as TIFF files

- PtGui for stitching.
  World leading commercial software but one of the risk factors as hasn't been applied at this scale.

- Developed custom tools to evaluate PtGui control point distribution.

- libvips library of image tools and to convert to various pyramidal formats for real time navigation.

- Developed custom image processing tools operating on the VIPS images
  - Lossless format,
  - No dimension limitations
  - Fast to process
  - Easy to develop additional tools.

PtGui

- Load initial grid of photographs

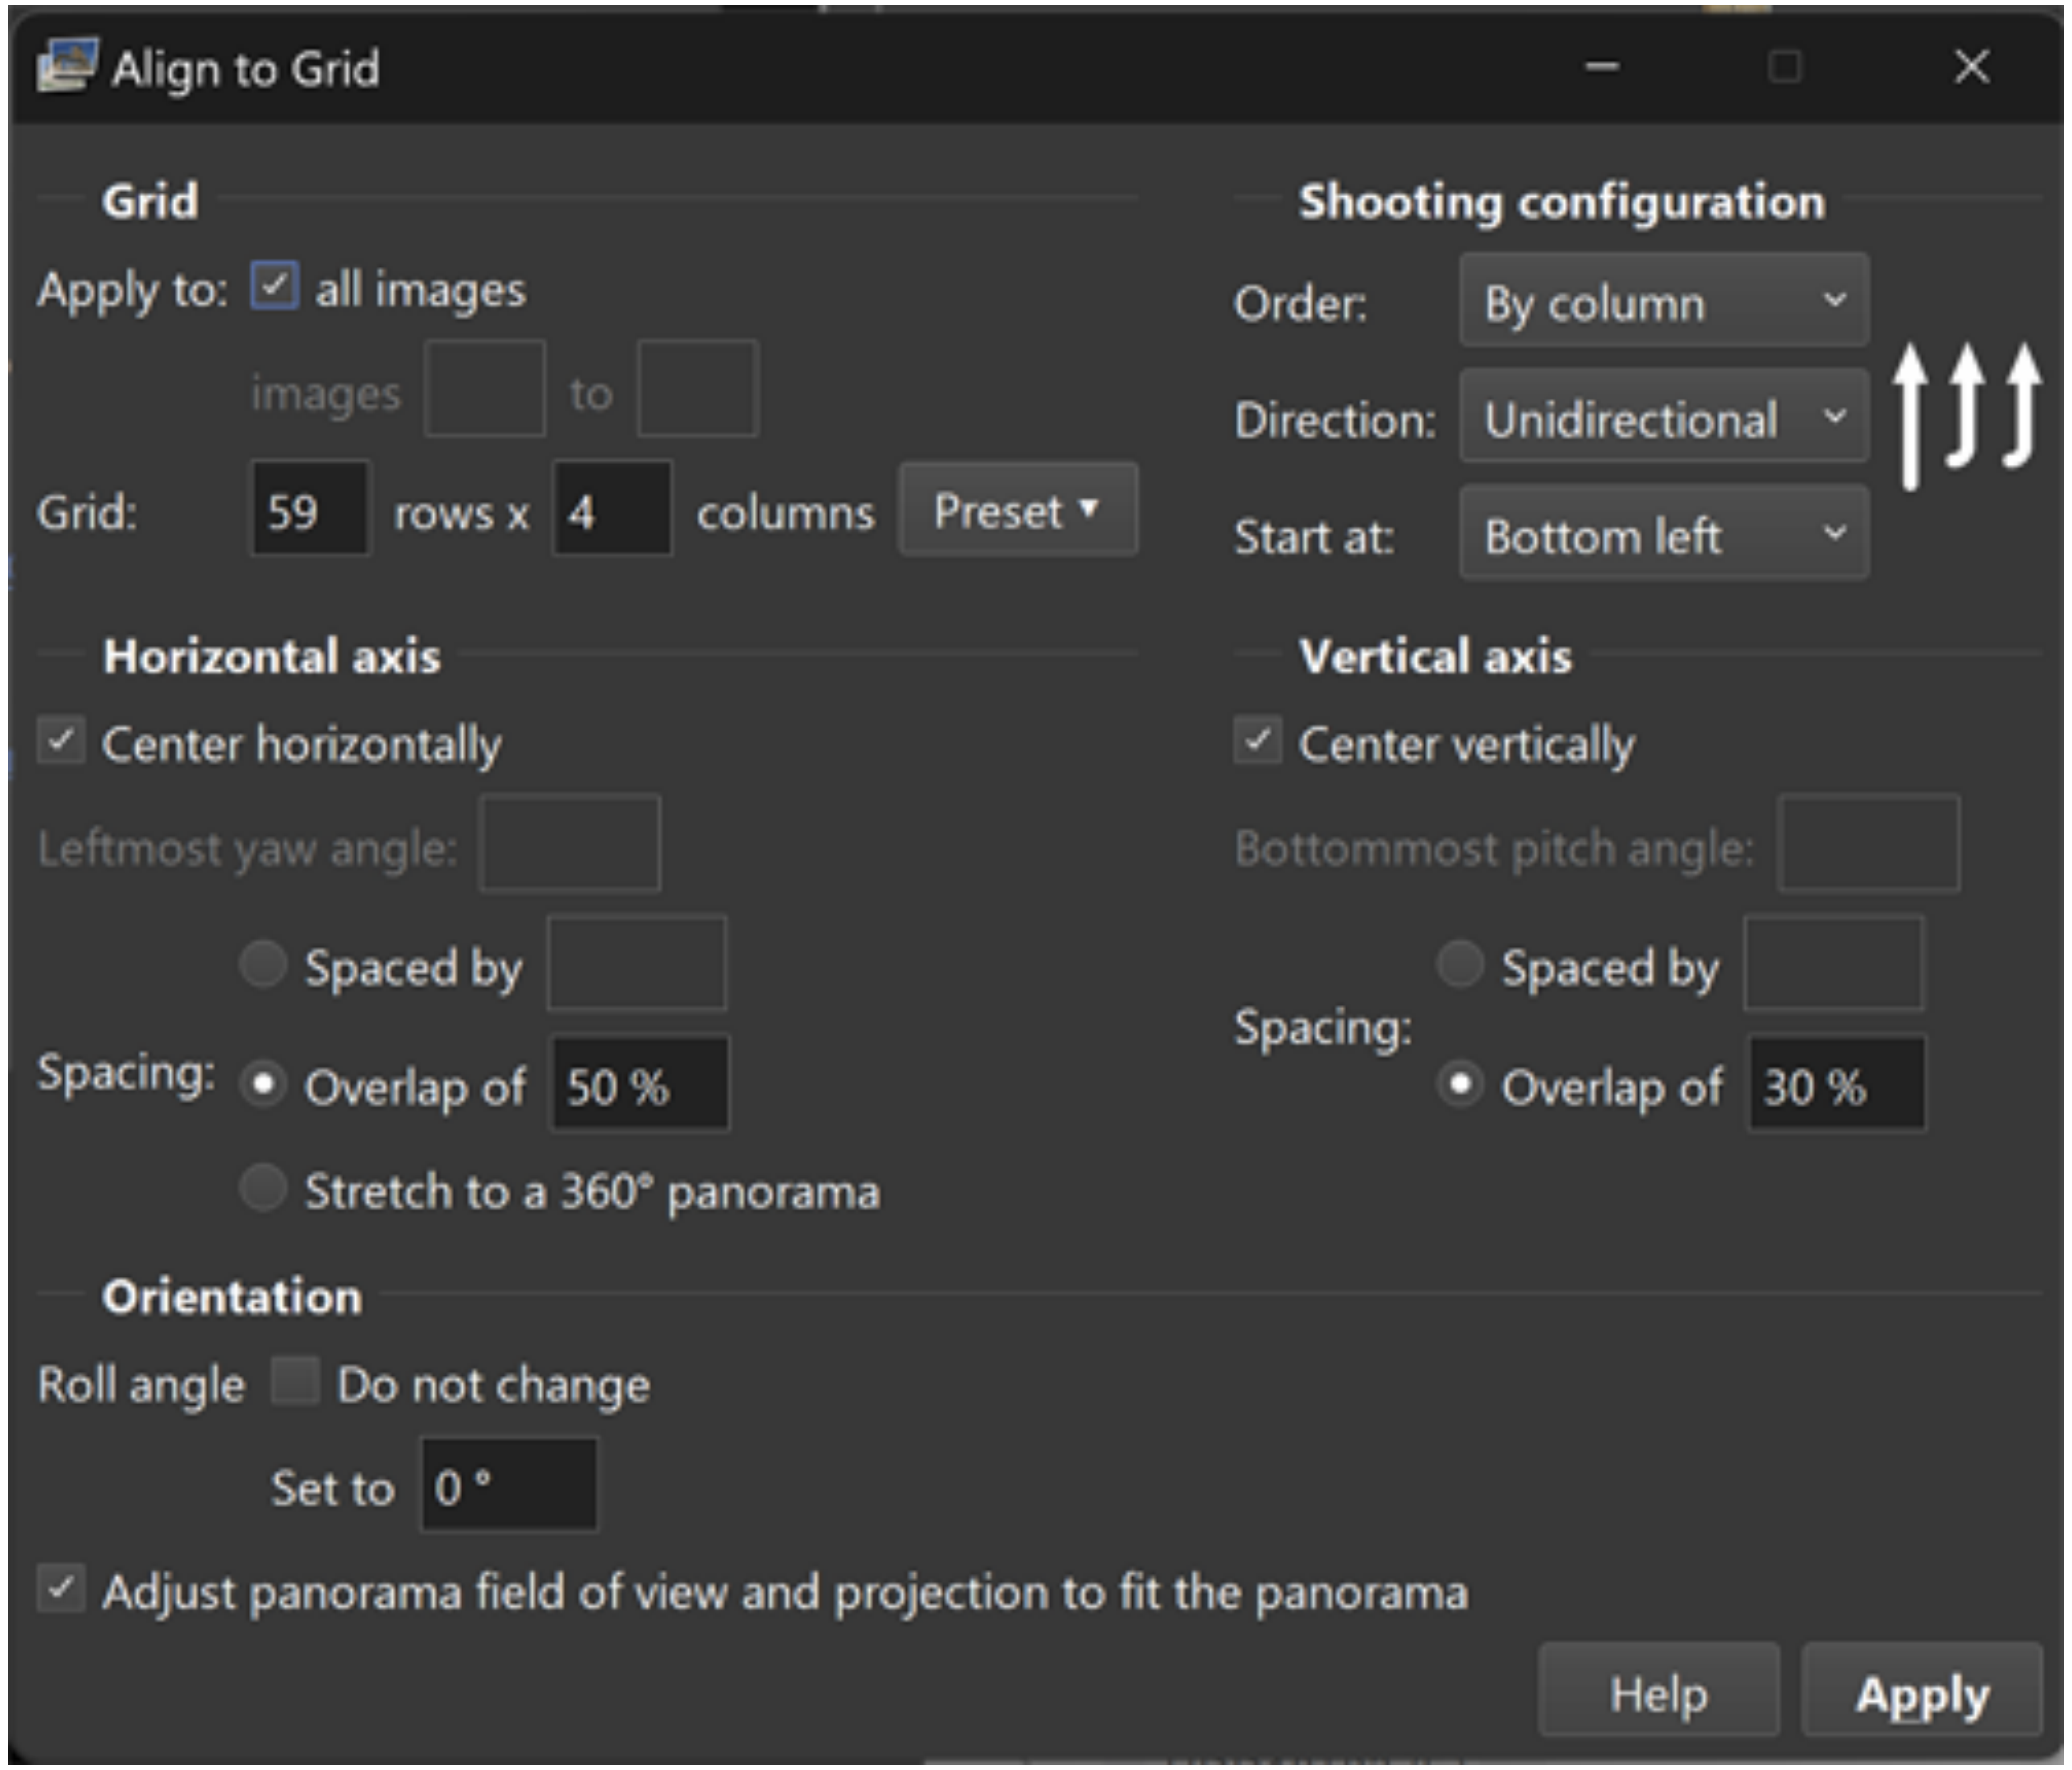

PtGui

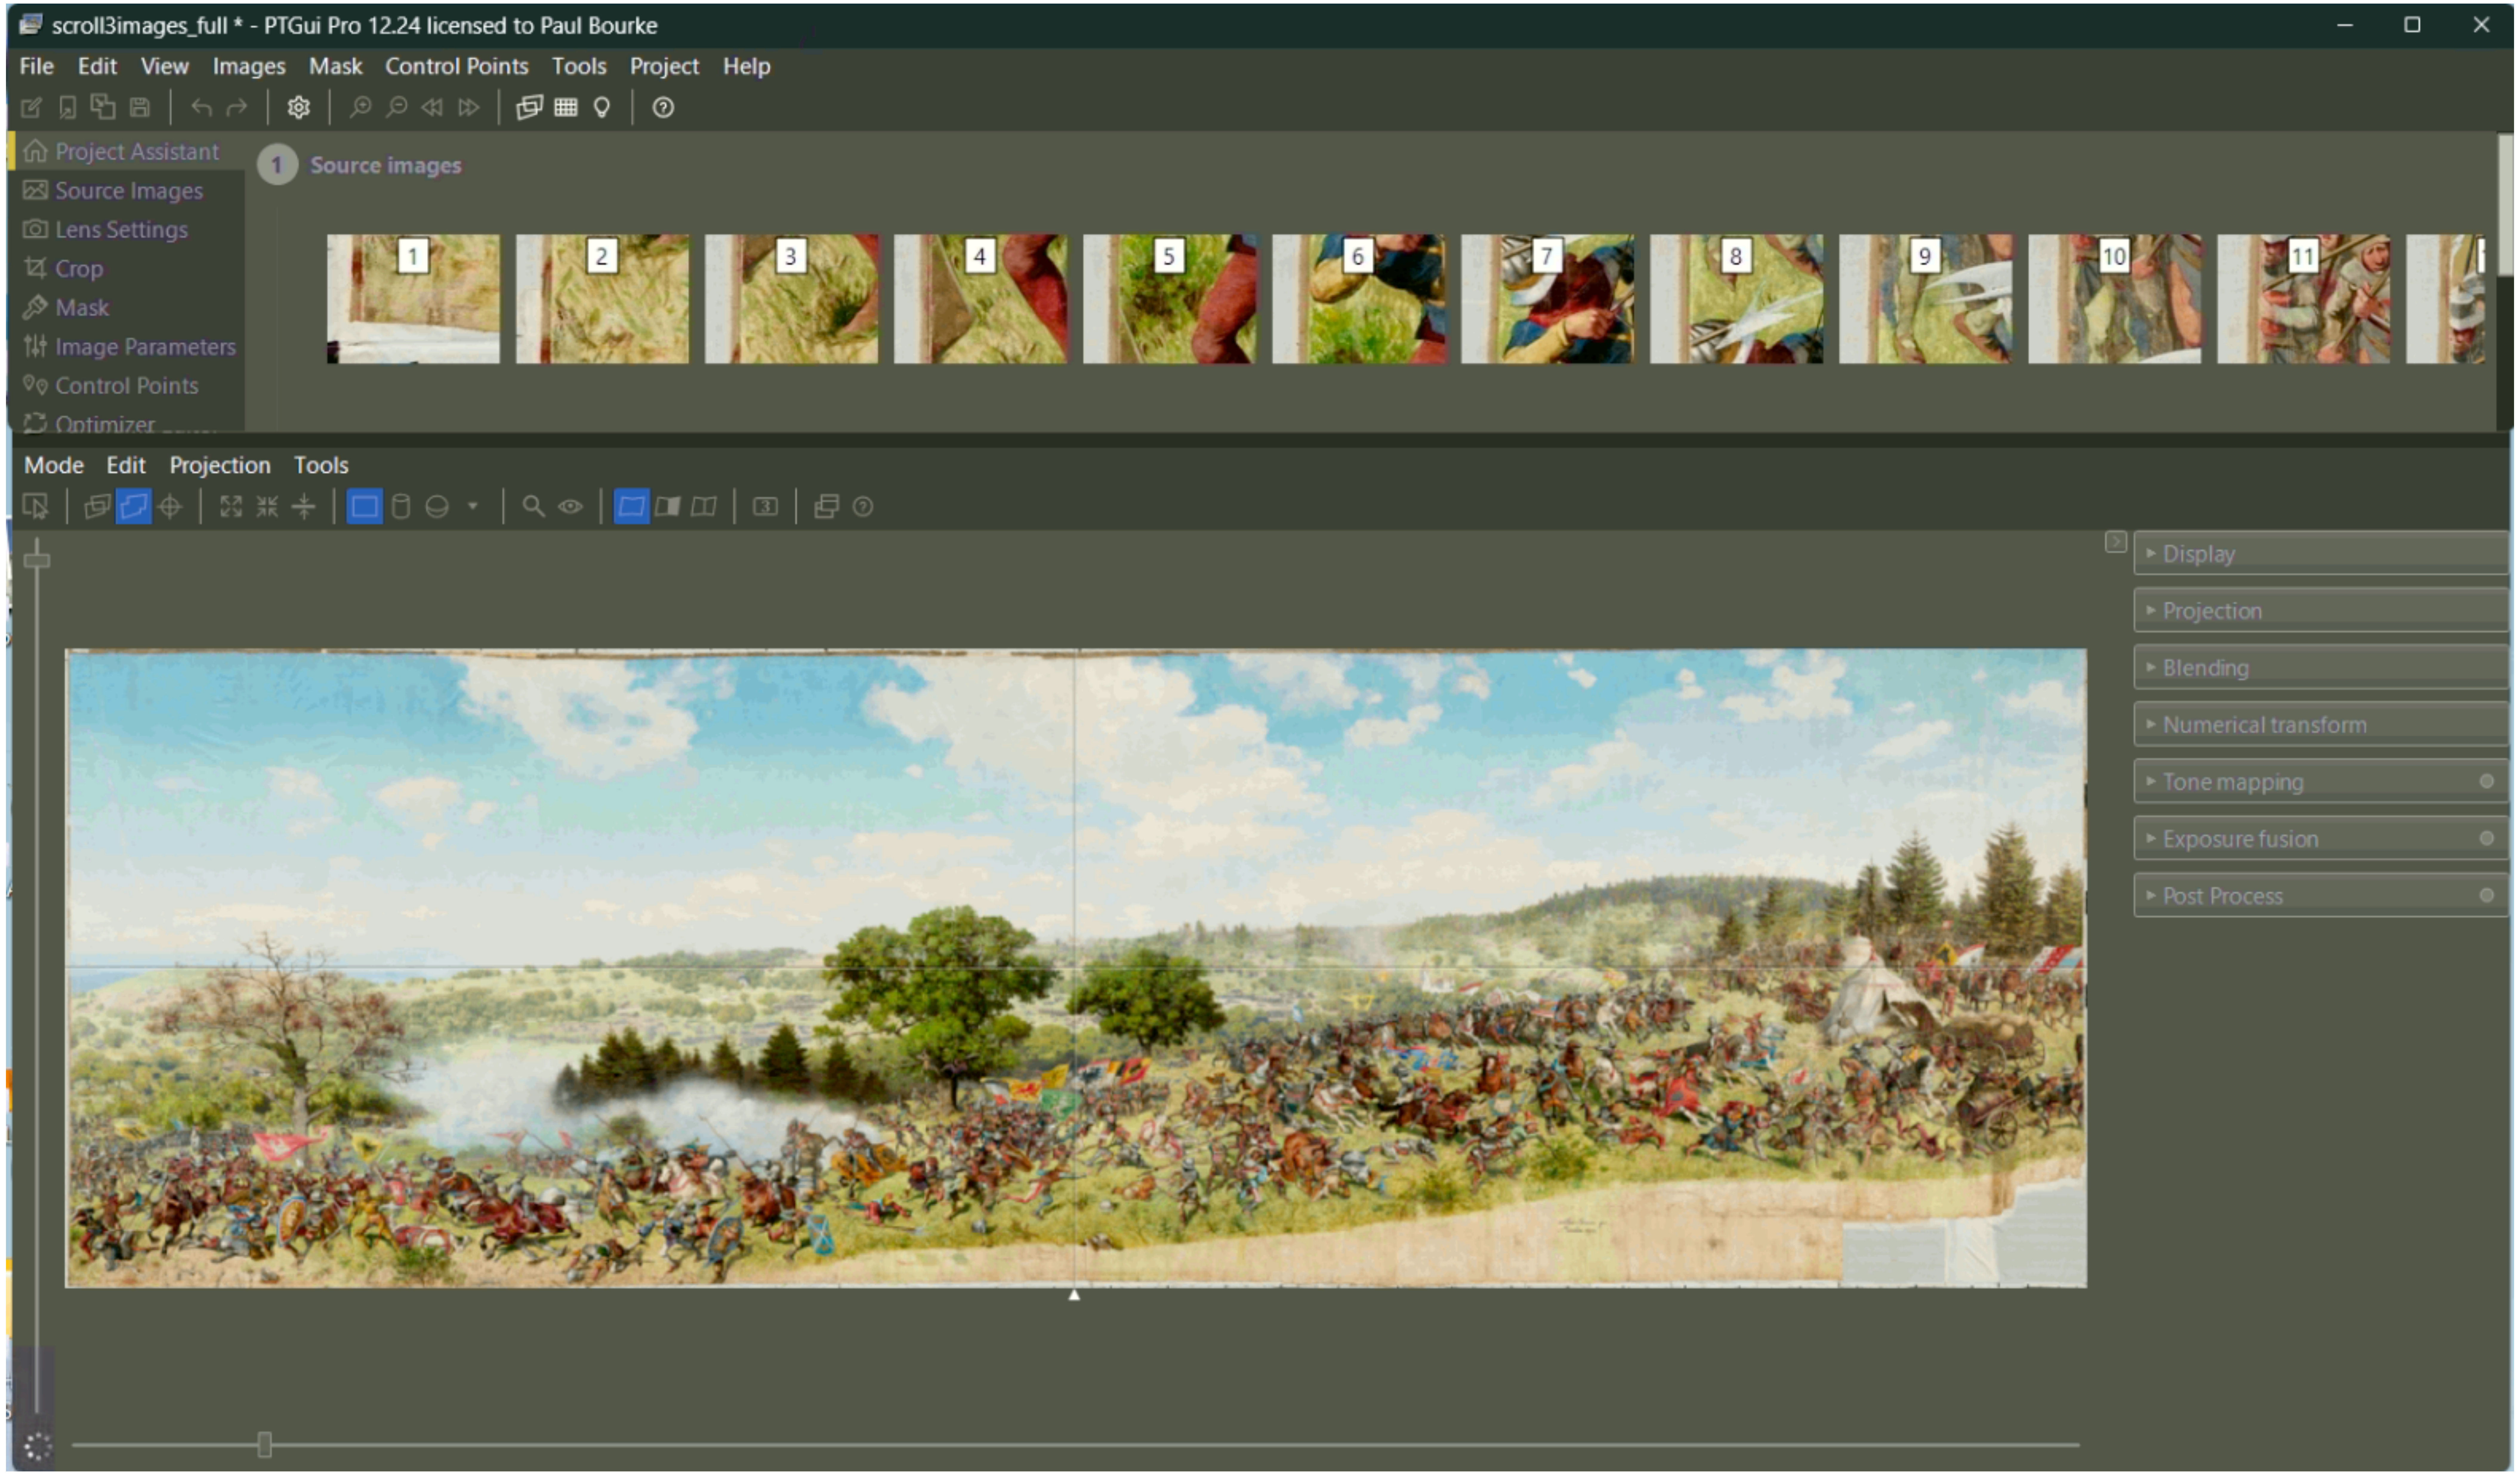
scroll3images_full * - PTGui Pro 12.24 licensed to Paul Bourke
File Edit View Images Mask Control Points Tools Project Help
Project Assistant
Source Images
Lens Settings
Crop
Mask
Image Parameters
Control Points
Optimizer
Source images
Mode Edit Projection Tools
Display
Projection
Blending
Numerical transform
Tone mapping
Exposure fusion
Post Process

Visualisation tool

- Shows images sharing control points

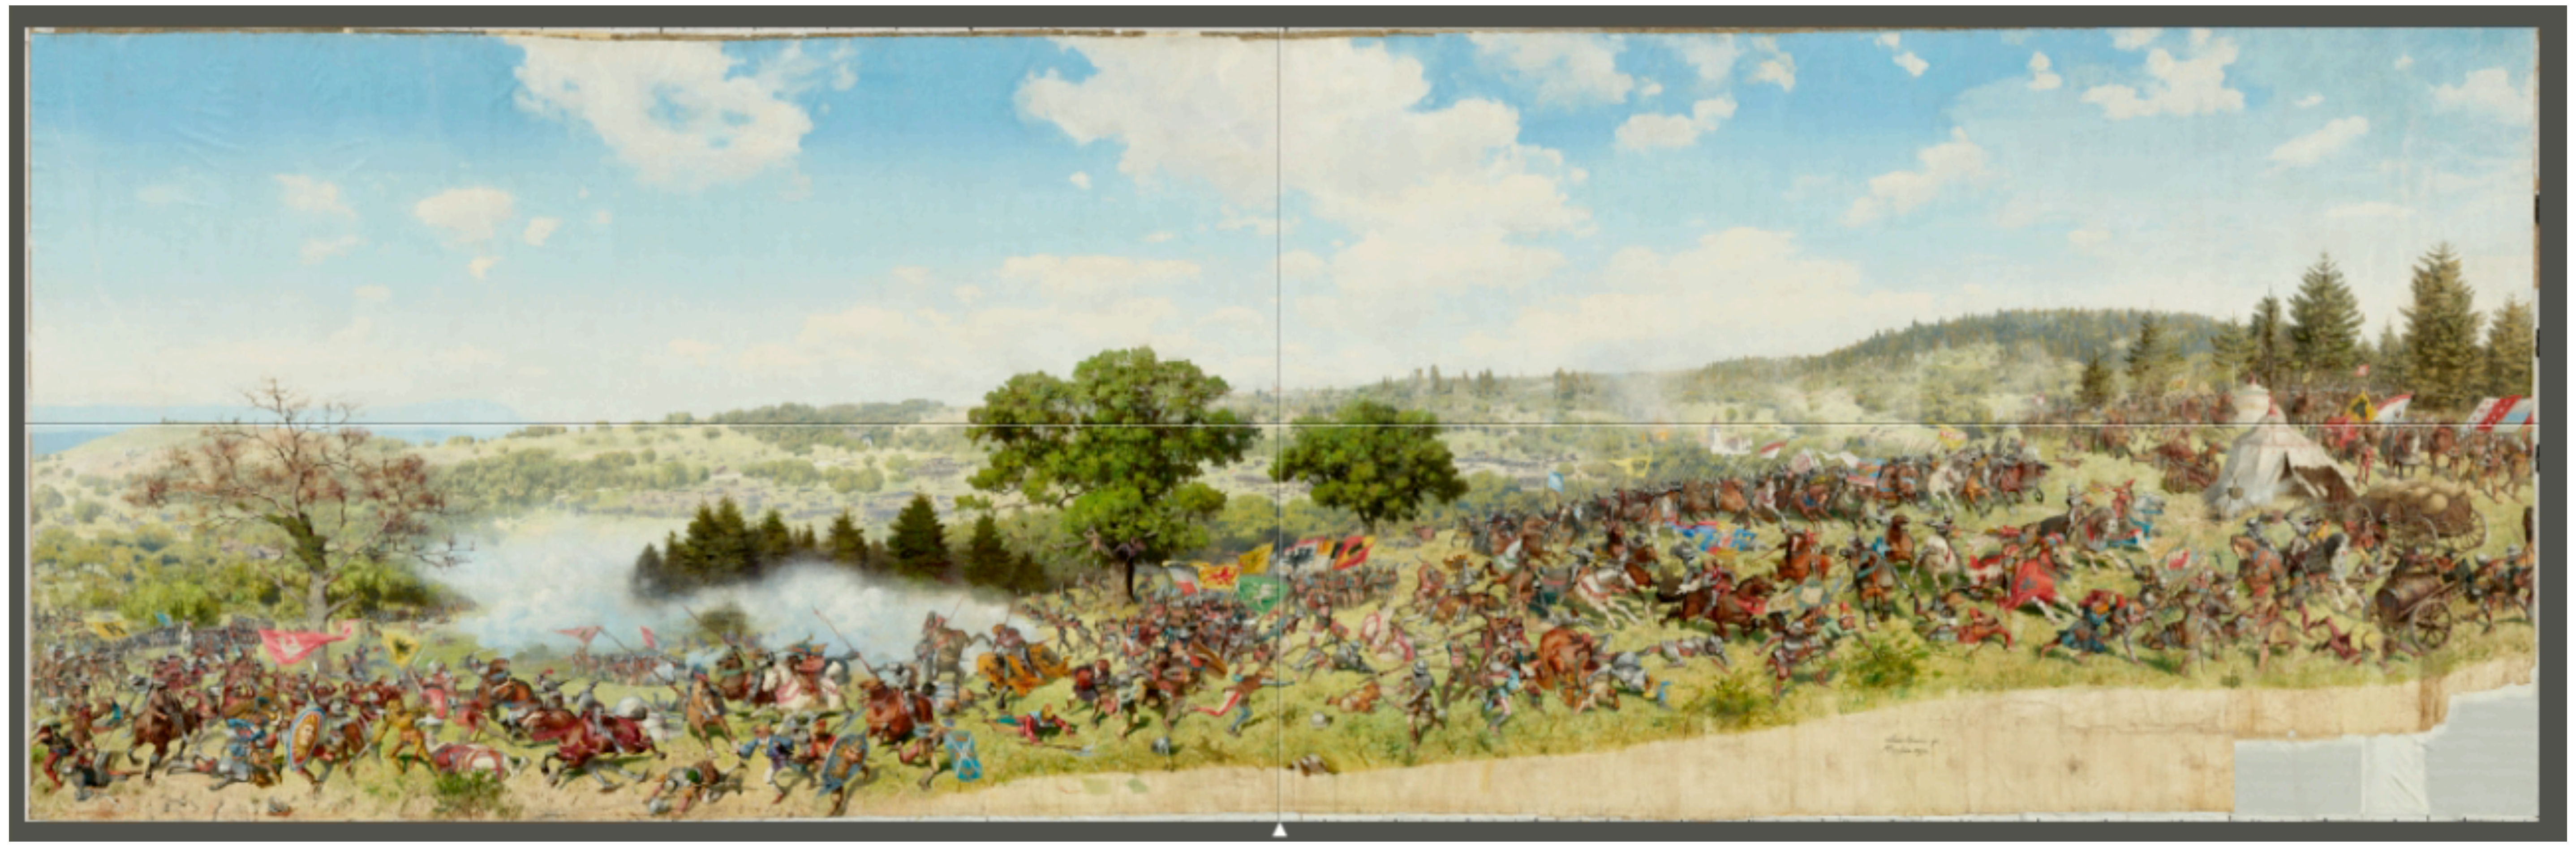

Forced control points

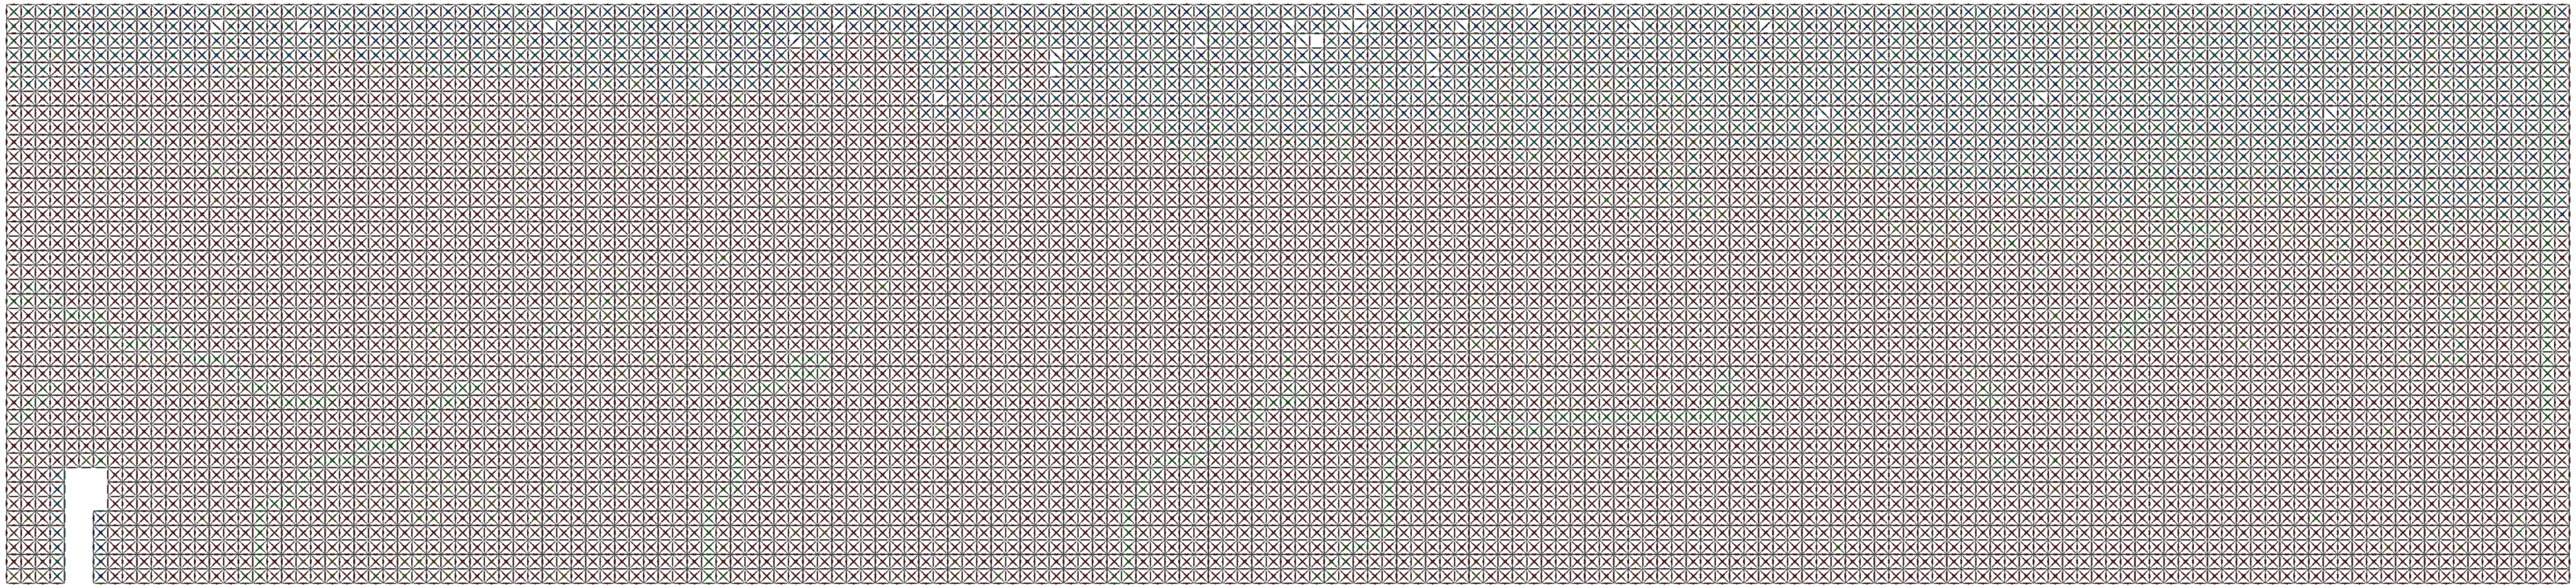

Image format and editing software?

- Most image formats we encounter on a regular basis (JPEG,PNG,TIFF) have size limitations, if not hard format limits, limits imposed by software.
- More sophisticated formats (BigTiff, JPEG2000) are complicated and have variable software support.
- We require a lossless format for the processing pipeline.
- Traditional image editing packages expect to load the image into RAM.
  This equates to the Murten image needing 4.8 TBytes of RAM!
- PhotoShop has a limit of a 300000x300000 pixel rectangle (90GB Ram).
  But even if one opened such an image, PhotoShop would be intolerably slow on any current hardware.

Red rectangle is the largest piece PhotoShop can potentially edit: 300,000 x 300,000 pixels

VIPS

- Consists of a image format and a library of image tools (libvips).
  Limited to dimensions up to 2^{32} pixels which is more than 4294 million pixels horizontally or vertically.
- In the end libvips, despite being designed for low memory usage on large images, failed for the image sizes required.
- Most problems due to it simply not having been tested on such large images, developers fixed some the issues encountered.
- Could not rely on bugs fixed in time so in the end I wrote all the image manipulation tools required.
- Not too difficult because the VIPS format is extremely simple, at least for the minimal number of features we required.
- It consists of a short 64 bytes header with information on dimensions, format layout, pixel depth …
- Followed by essentially raw r,g,b data. Uncompressed but this means very fast to read.
  Noting that lossless compression for high quality images doesn't achieve much.
  => A clever compressed format would not give any advantages in storage and would be slow.

Hyperboloid (again)

- The panorama cannot be represented as a 2D image without distortion
- This is exactly the same problem as trying to represent a Earth map on a flat plane
- If the painting was digitised at a uniform sampling (40 pixels/mm) there would need to be fewer pixels at the mid height than the top and bottom row

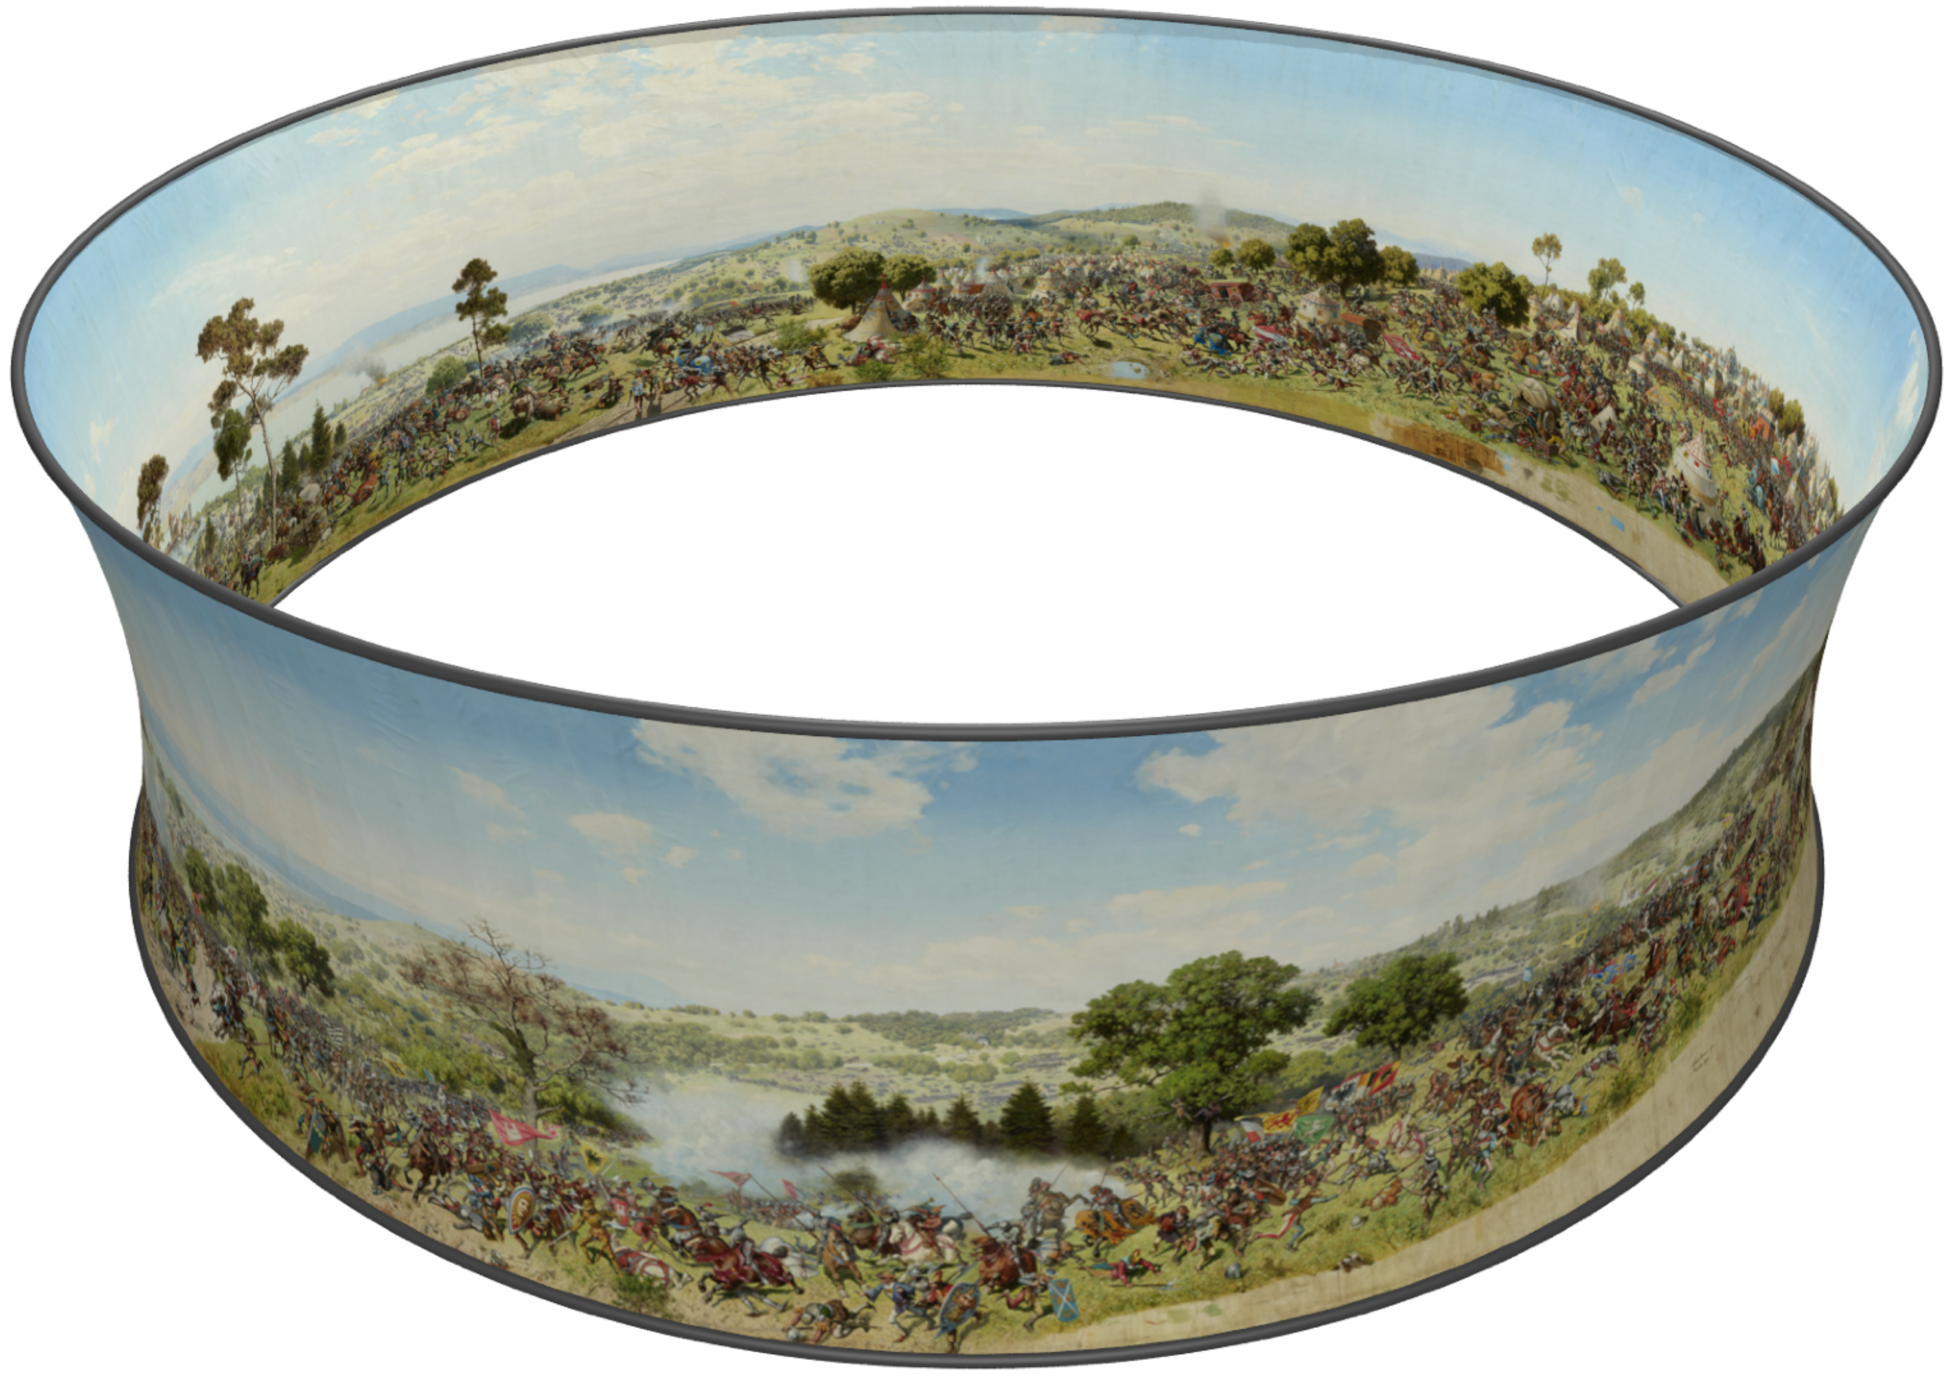

Hyperboloid correction

- Polynomial function fitted to each edge
- Image warped to form a straight edge at each end

Rolle Nr. I	Rollenlänge I gesamt
- *Oberkante* H^2= 0	3435,9 cm
- *Bahnhälfte* H= 510 cm	3348,5 cm
- *Unterkante*	3427,3 cm

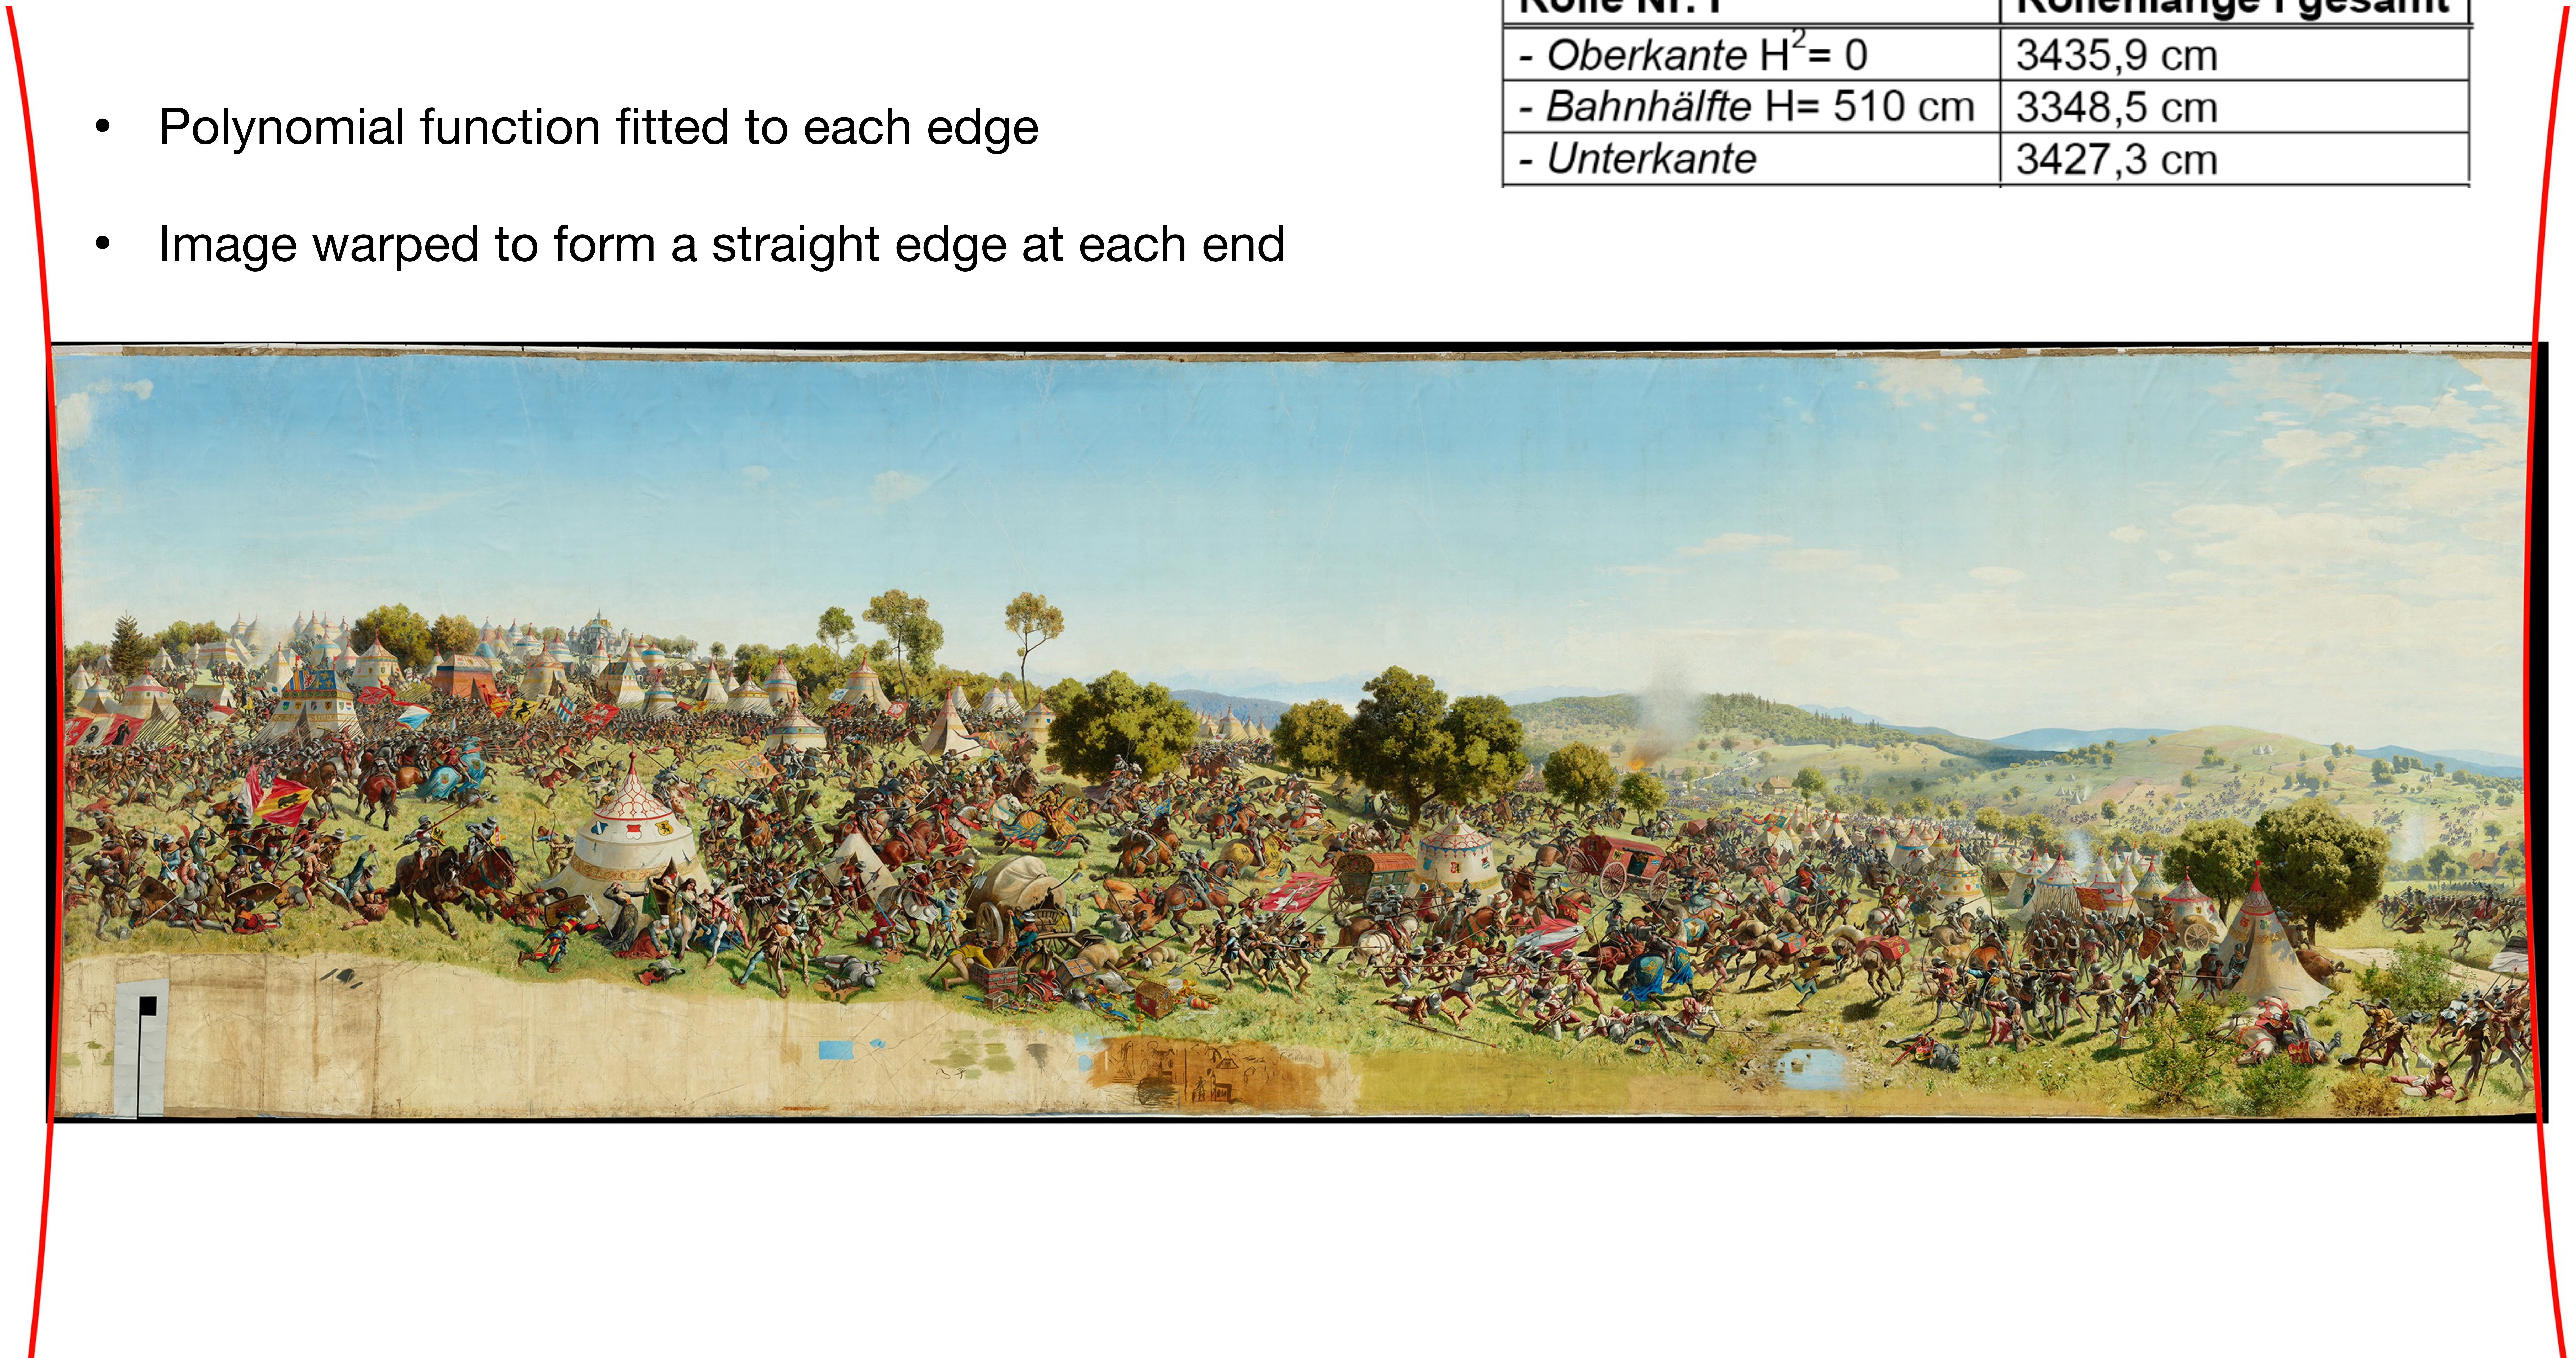

VIPS tool “vdecurve"

- Corrects the hyperboloid effect by stretching the rows by the appropriate amount. The degree of stretching is determined by fitting a fourth order polynomial to the left and right edge of the scroll image.

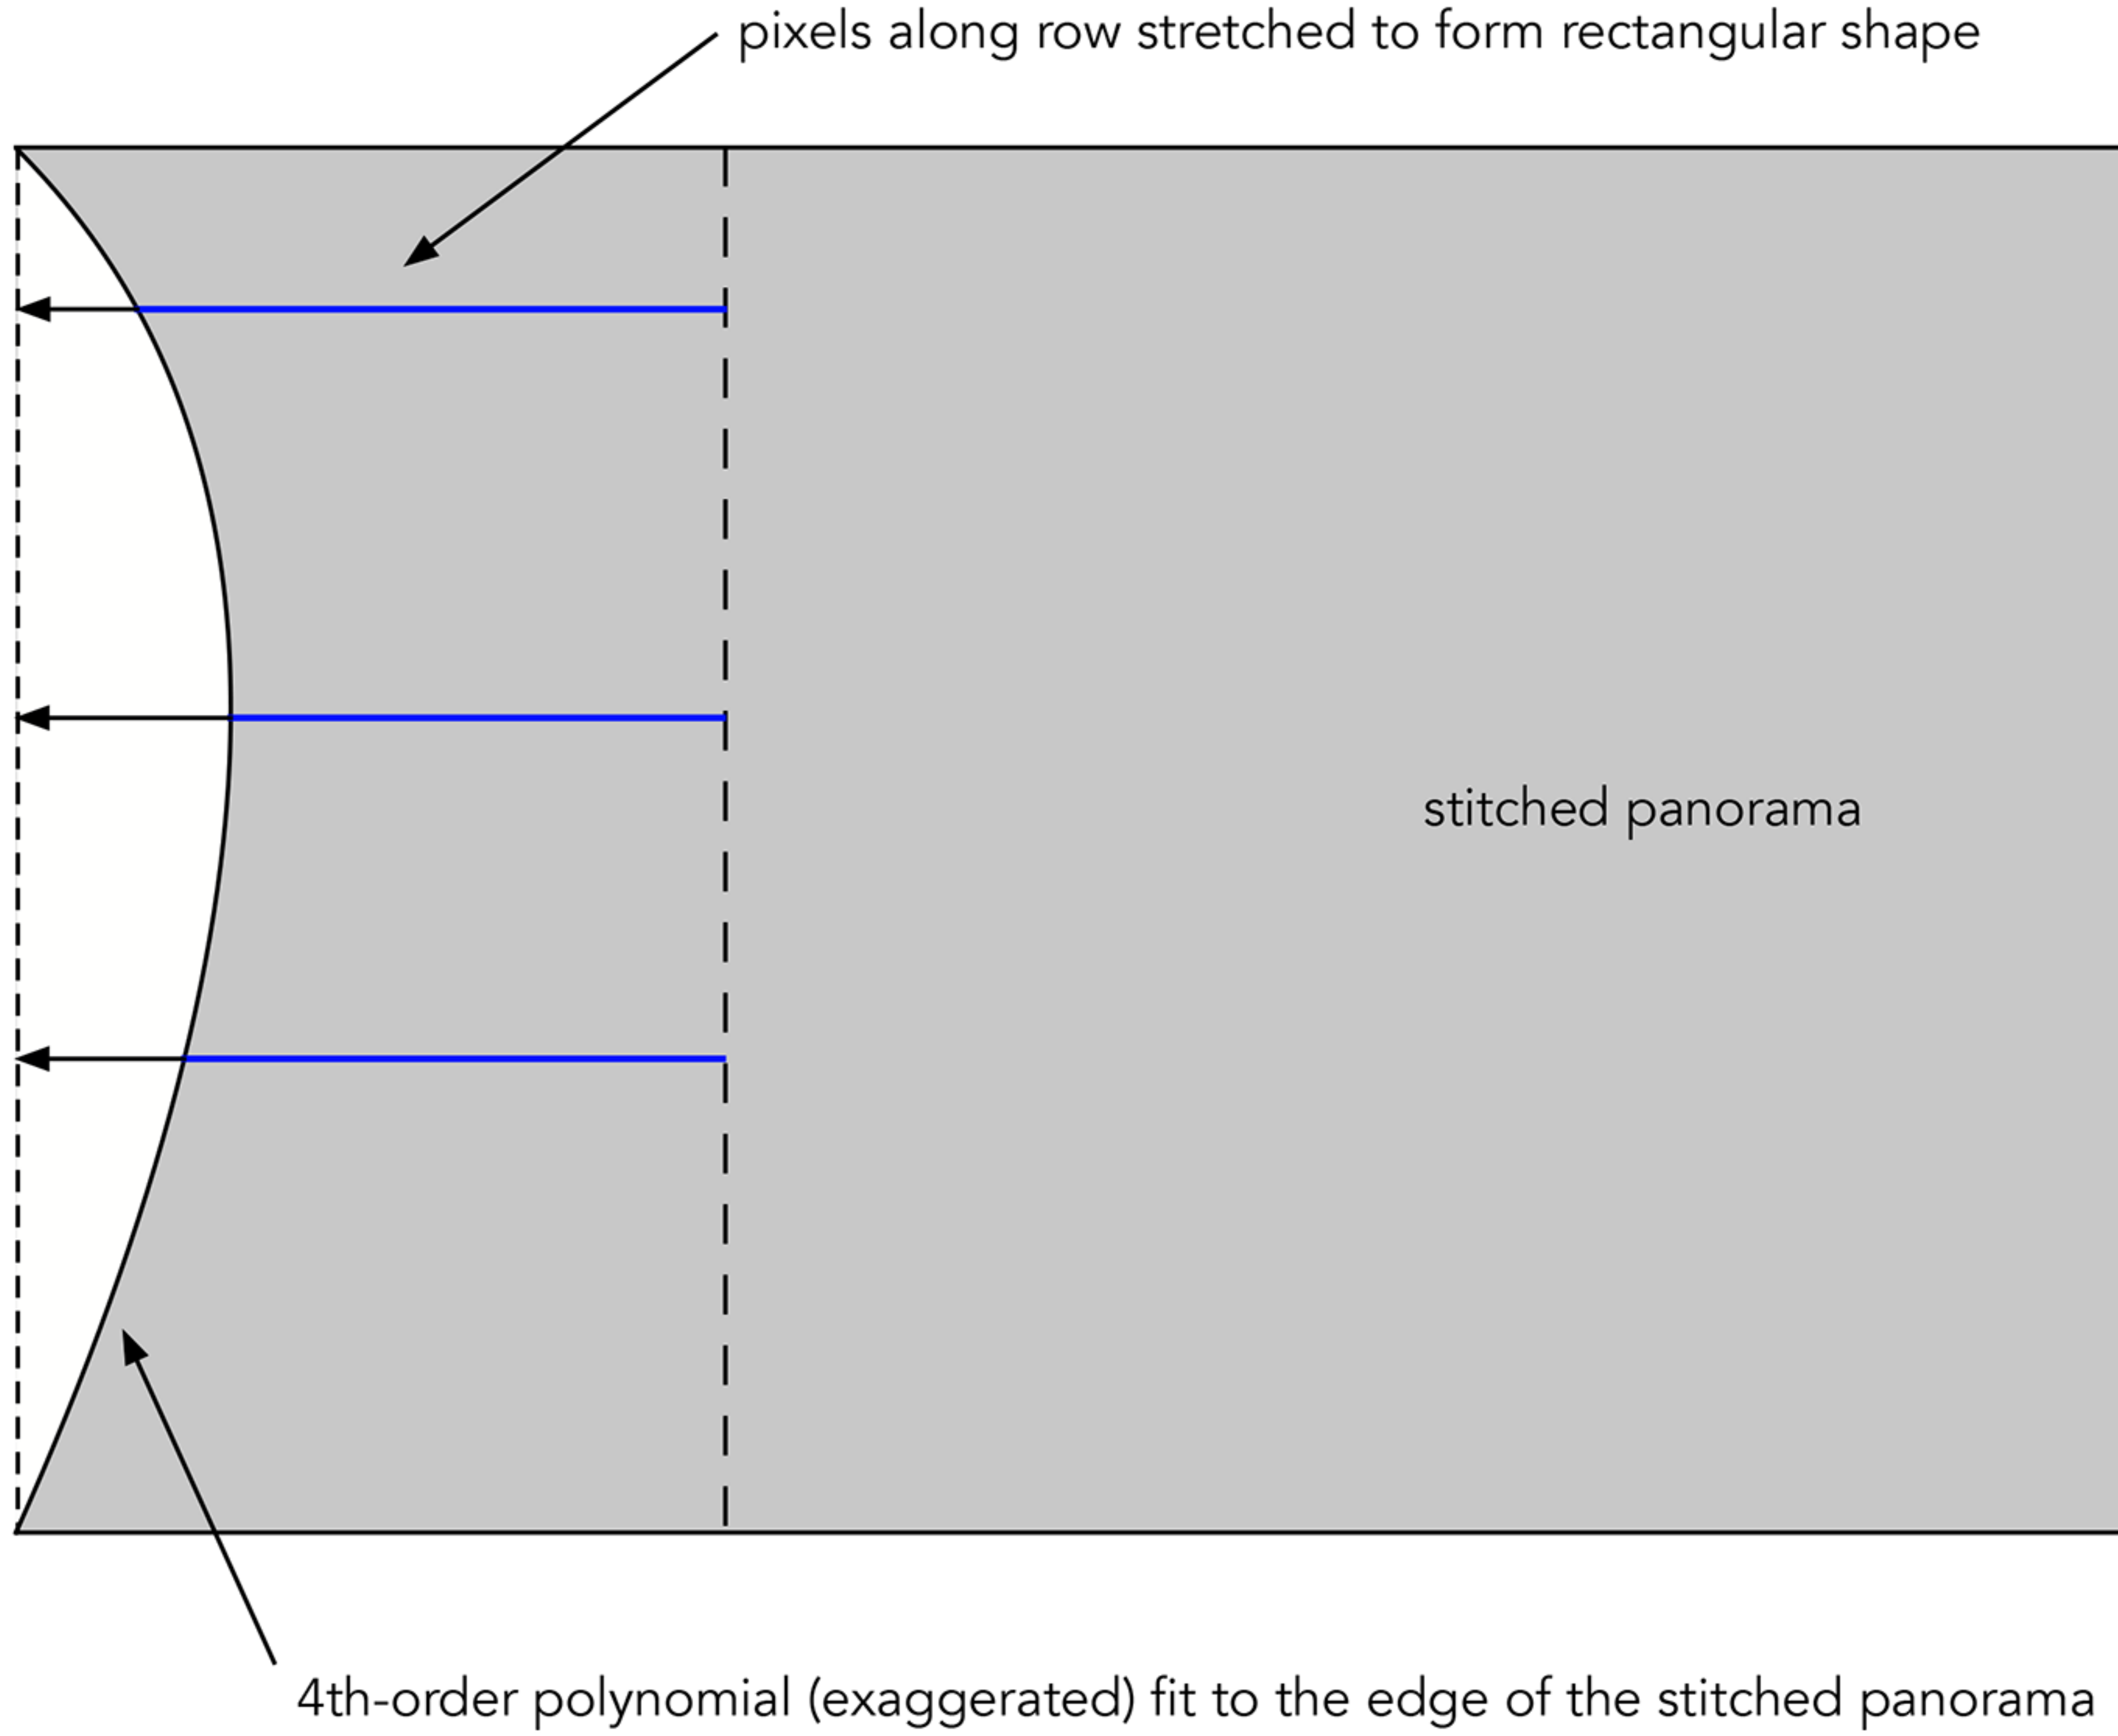

VIPS tool "valign"

- Allows each scroll to be vertically aligned with it's neighbour. The need for this arises from the jagged nature of the edges of the scroll and that they were not necessarily mounted perfectly straight on the roll. Corresponding points are identified along adjacent edges and a linear scale mapping performed to ensure a match.

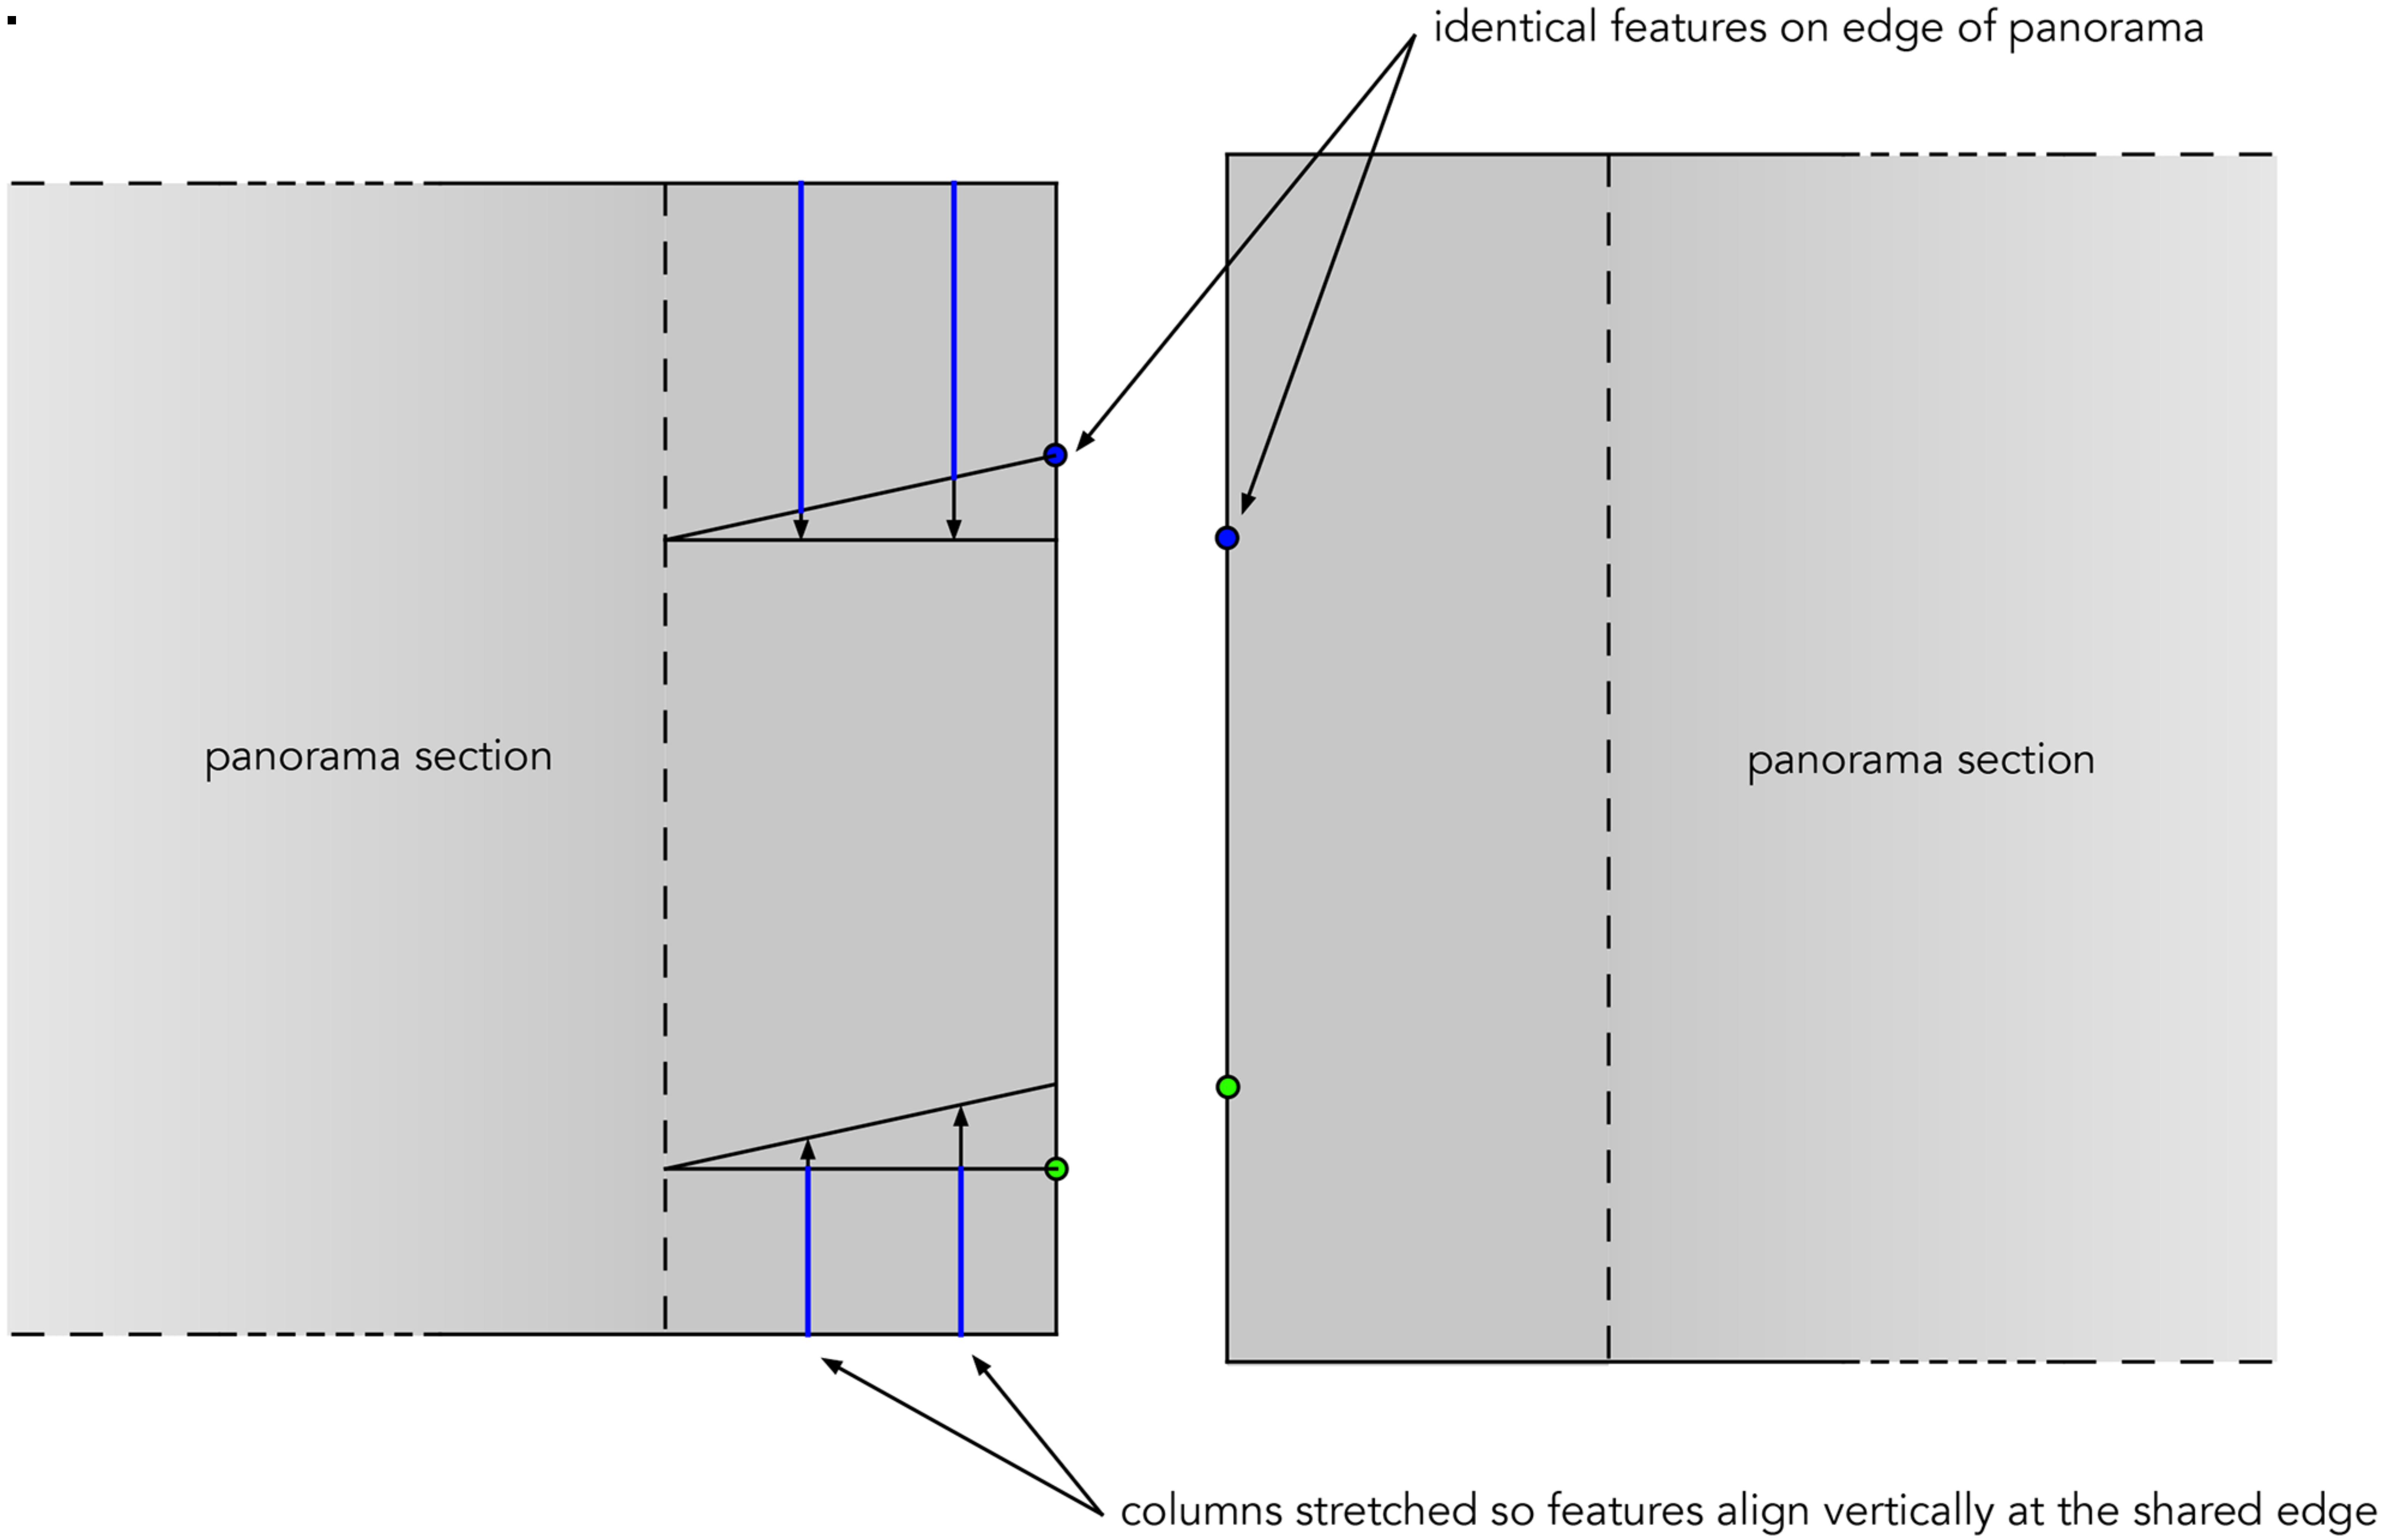

VIPS tools

- **vsplit** and **vjoin**: Splits an image into pieces (vertically or horizontally) and joins them back together. Primarily used to create smaller work items for processing tasks that only required a portion of a scroll.
- **vextract** and **vinsert**: Extract an arbitrary chunk of the painting, possibly perform some image editing, and reconstitute the edited chunk into the original. In order to avoid multiple copies this reconstitution was performed in place.
- **vcrop**: Crop an image, mainly required to remove the over-scanned boundary pieces.
- **vresize**: Subsample an image by some factor, primarily used to create smaller assets and create lower resolution versions for previewing/checking.
- **raw2v**: In those cases where PhotoShop was used to perform edits, a raw image was exported due to PhotoShop not supporting BigTiff.
- All of these image editing tools operated on the images in a row by row, or column by column order, as such only a single row or column was in memory at ay time. Feasible because images were stored on fast SSD drives.

Pipeline

- Individual images exported from CaptureOne DH

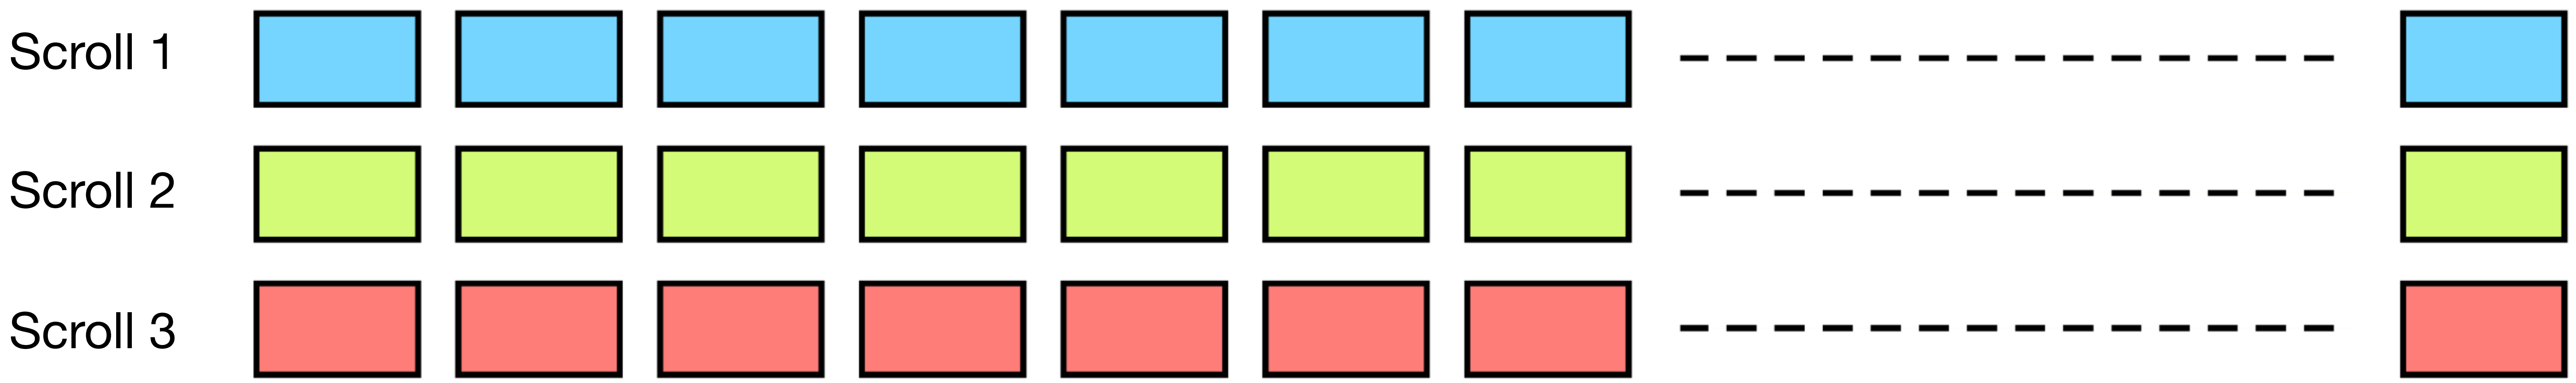

- Each scroll stitched with PtGui

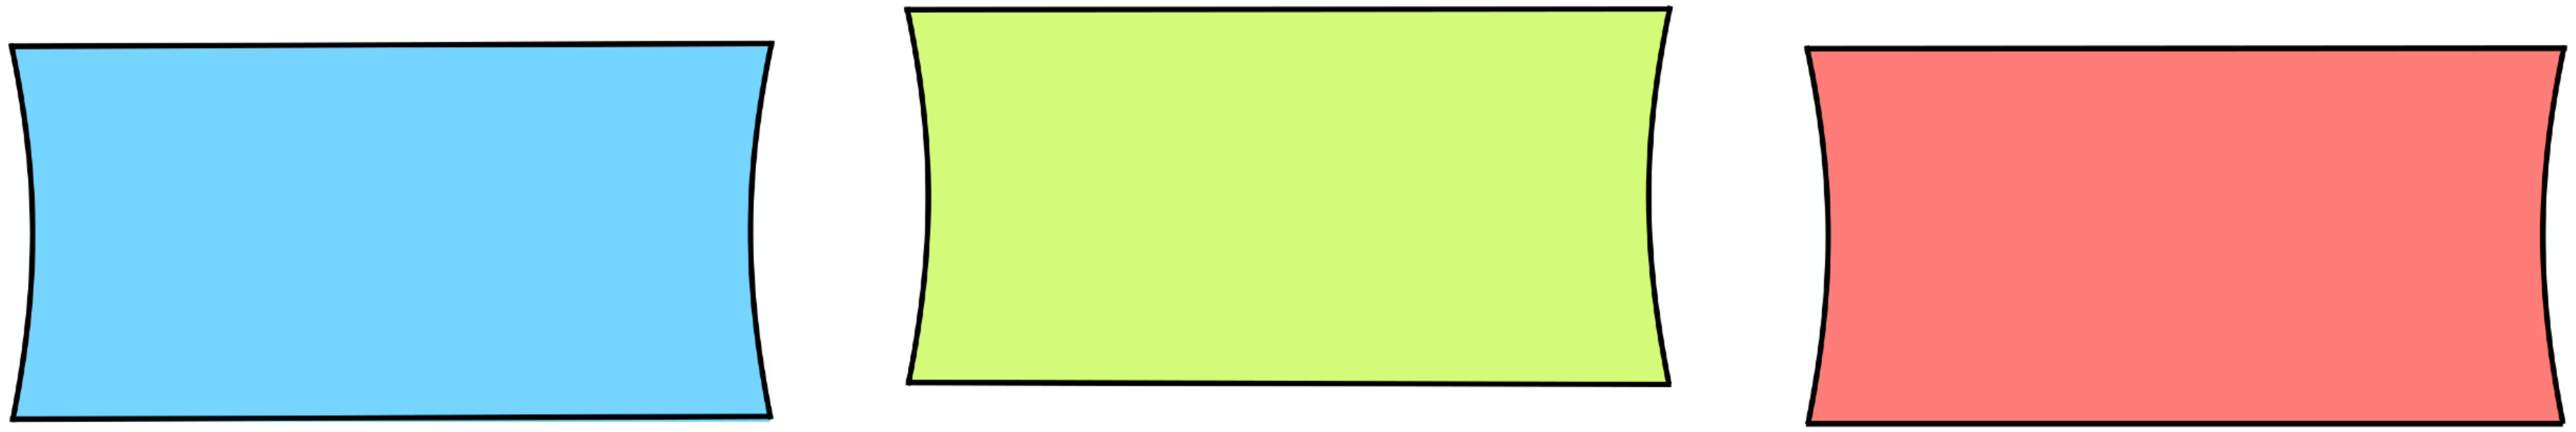

- Each scroll dehyperbolised with “vdecurve”

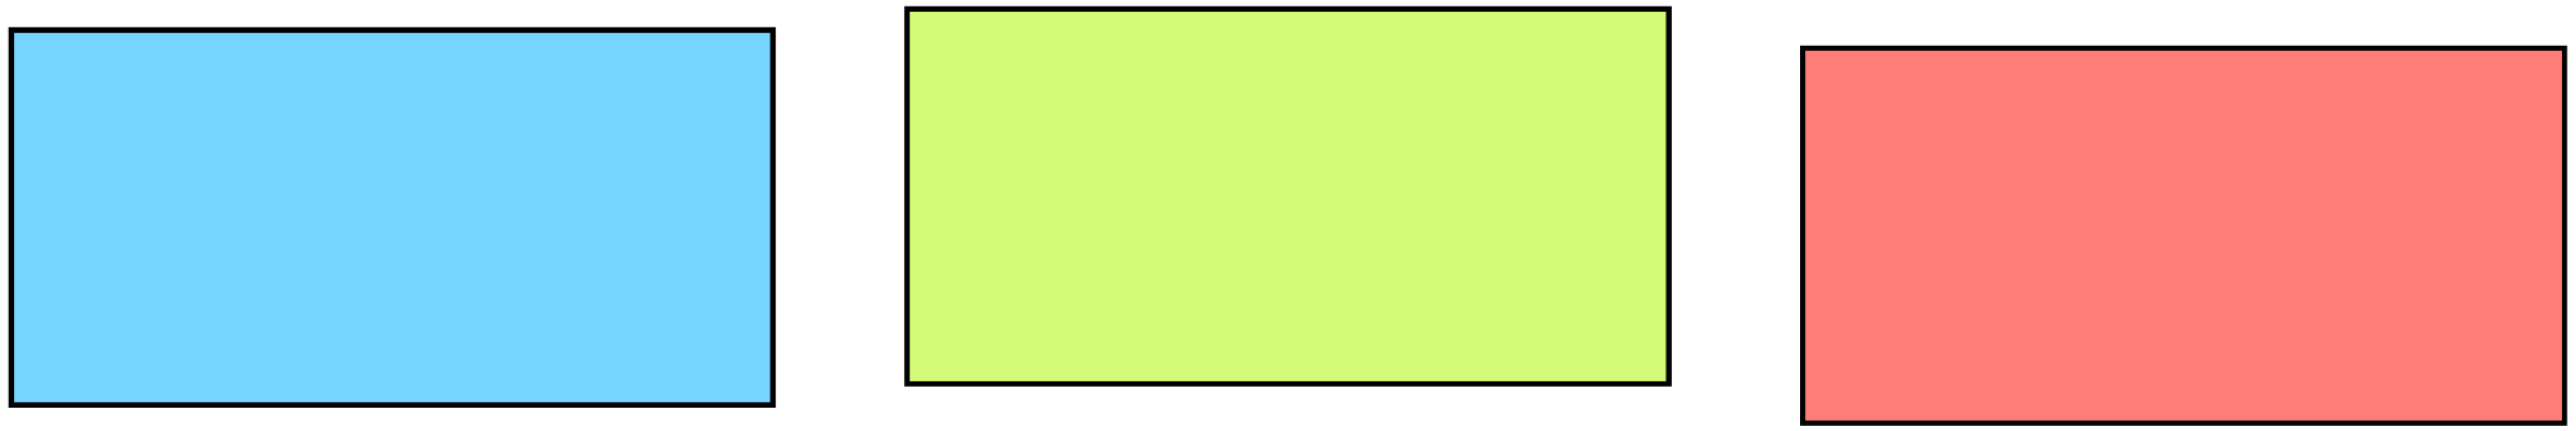

Pipeline

- Each scroll vertically aligned with “valign”

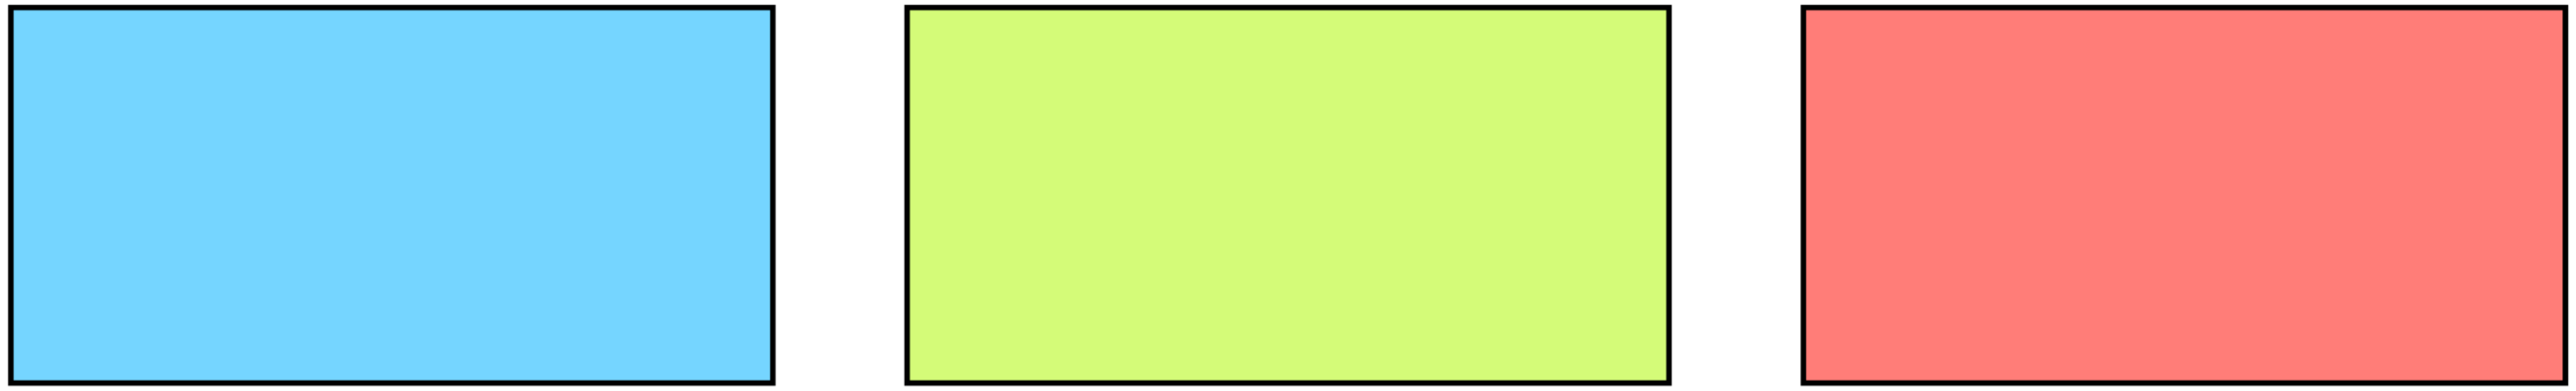

- Rearranged to facilitate wrapping and merged with “vsplit” and “vjoin”

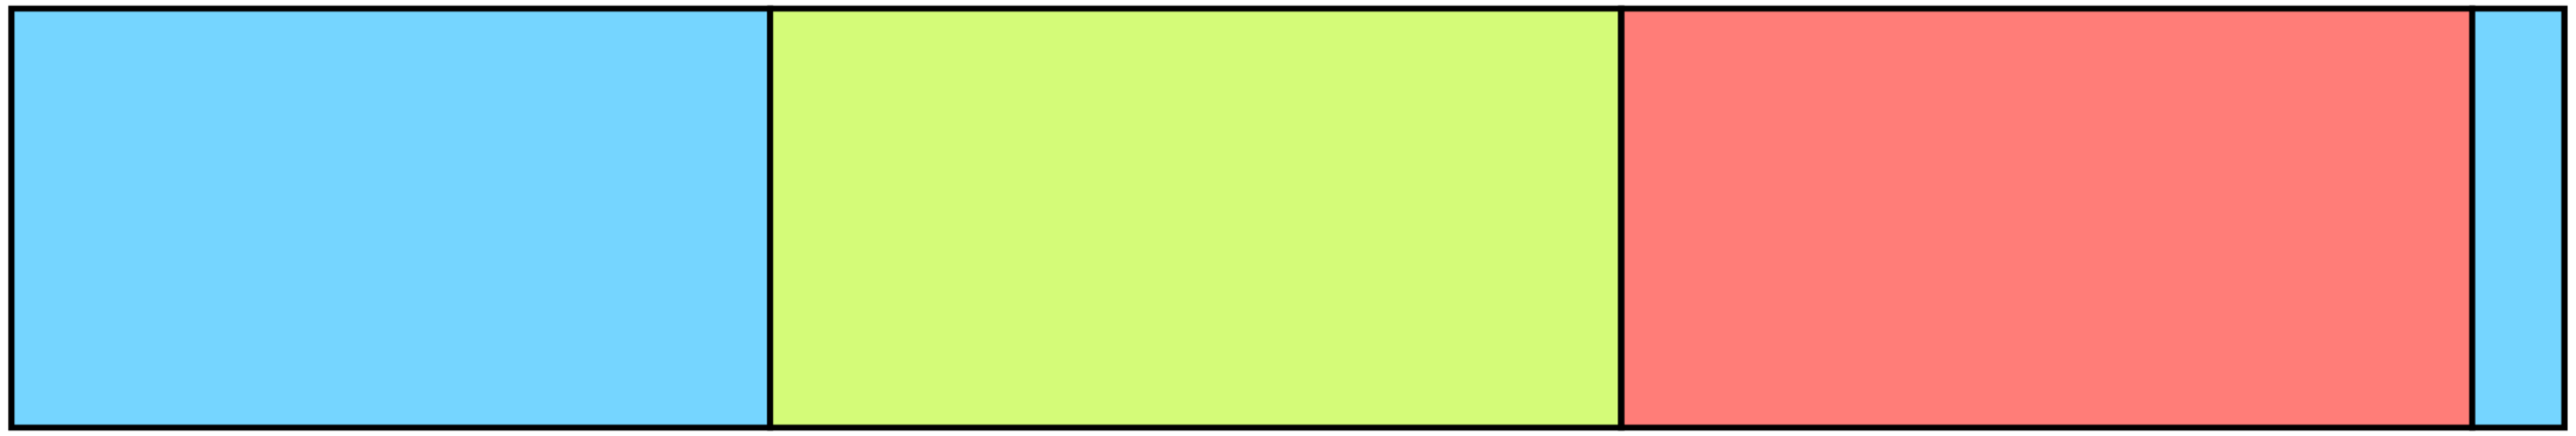

- Extract sky and ground patches, extend and reinsert with “vextract” and “vinsert”

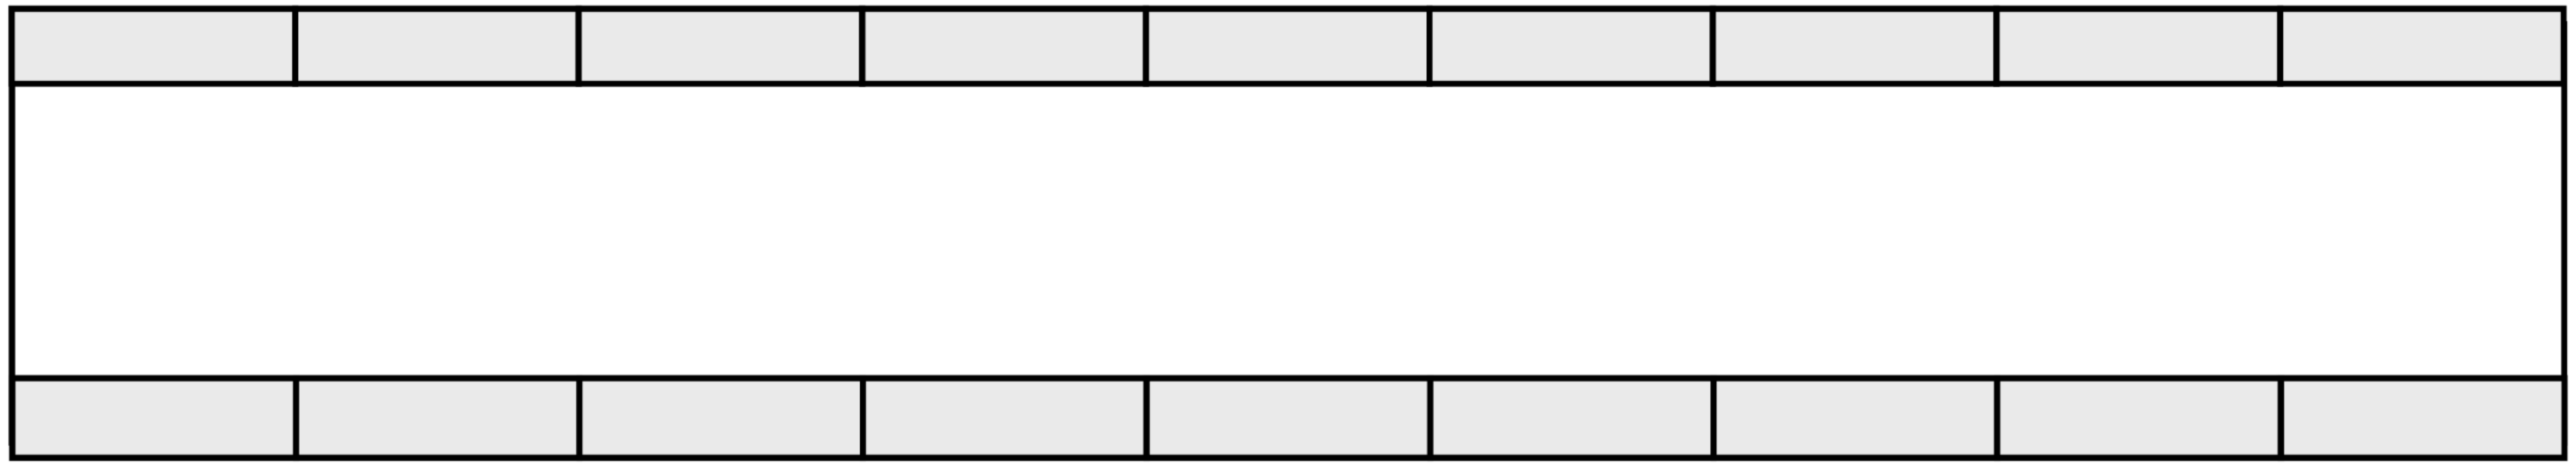

Pipeline

- Extract unpainted areas, edit and reinsert with “vextract” and “vinsert”

- Extract seams, edit and reinsert with “vextract” and “vinsert”

- Crop top and bottom with “vcrop”

Scanning hardware configuration

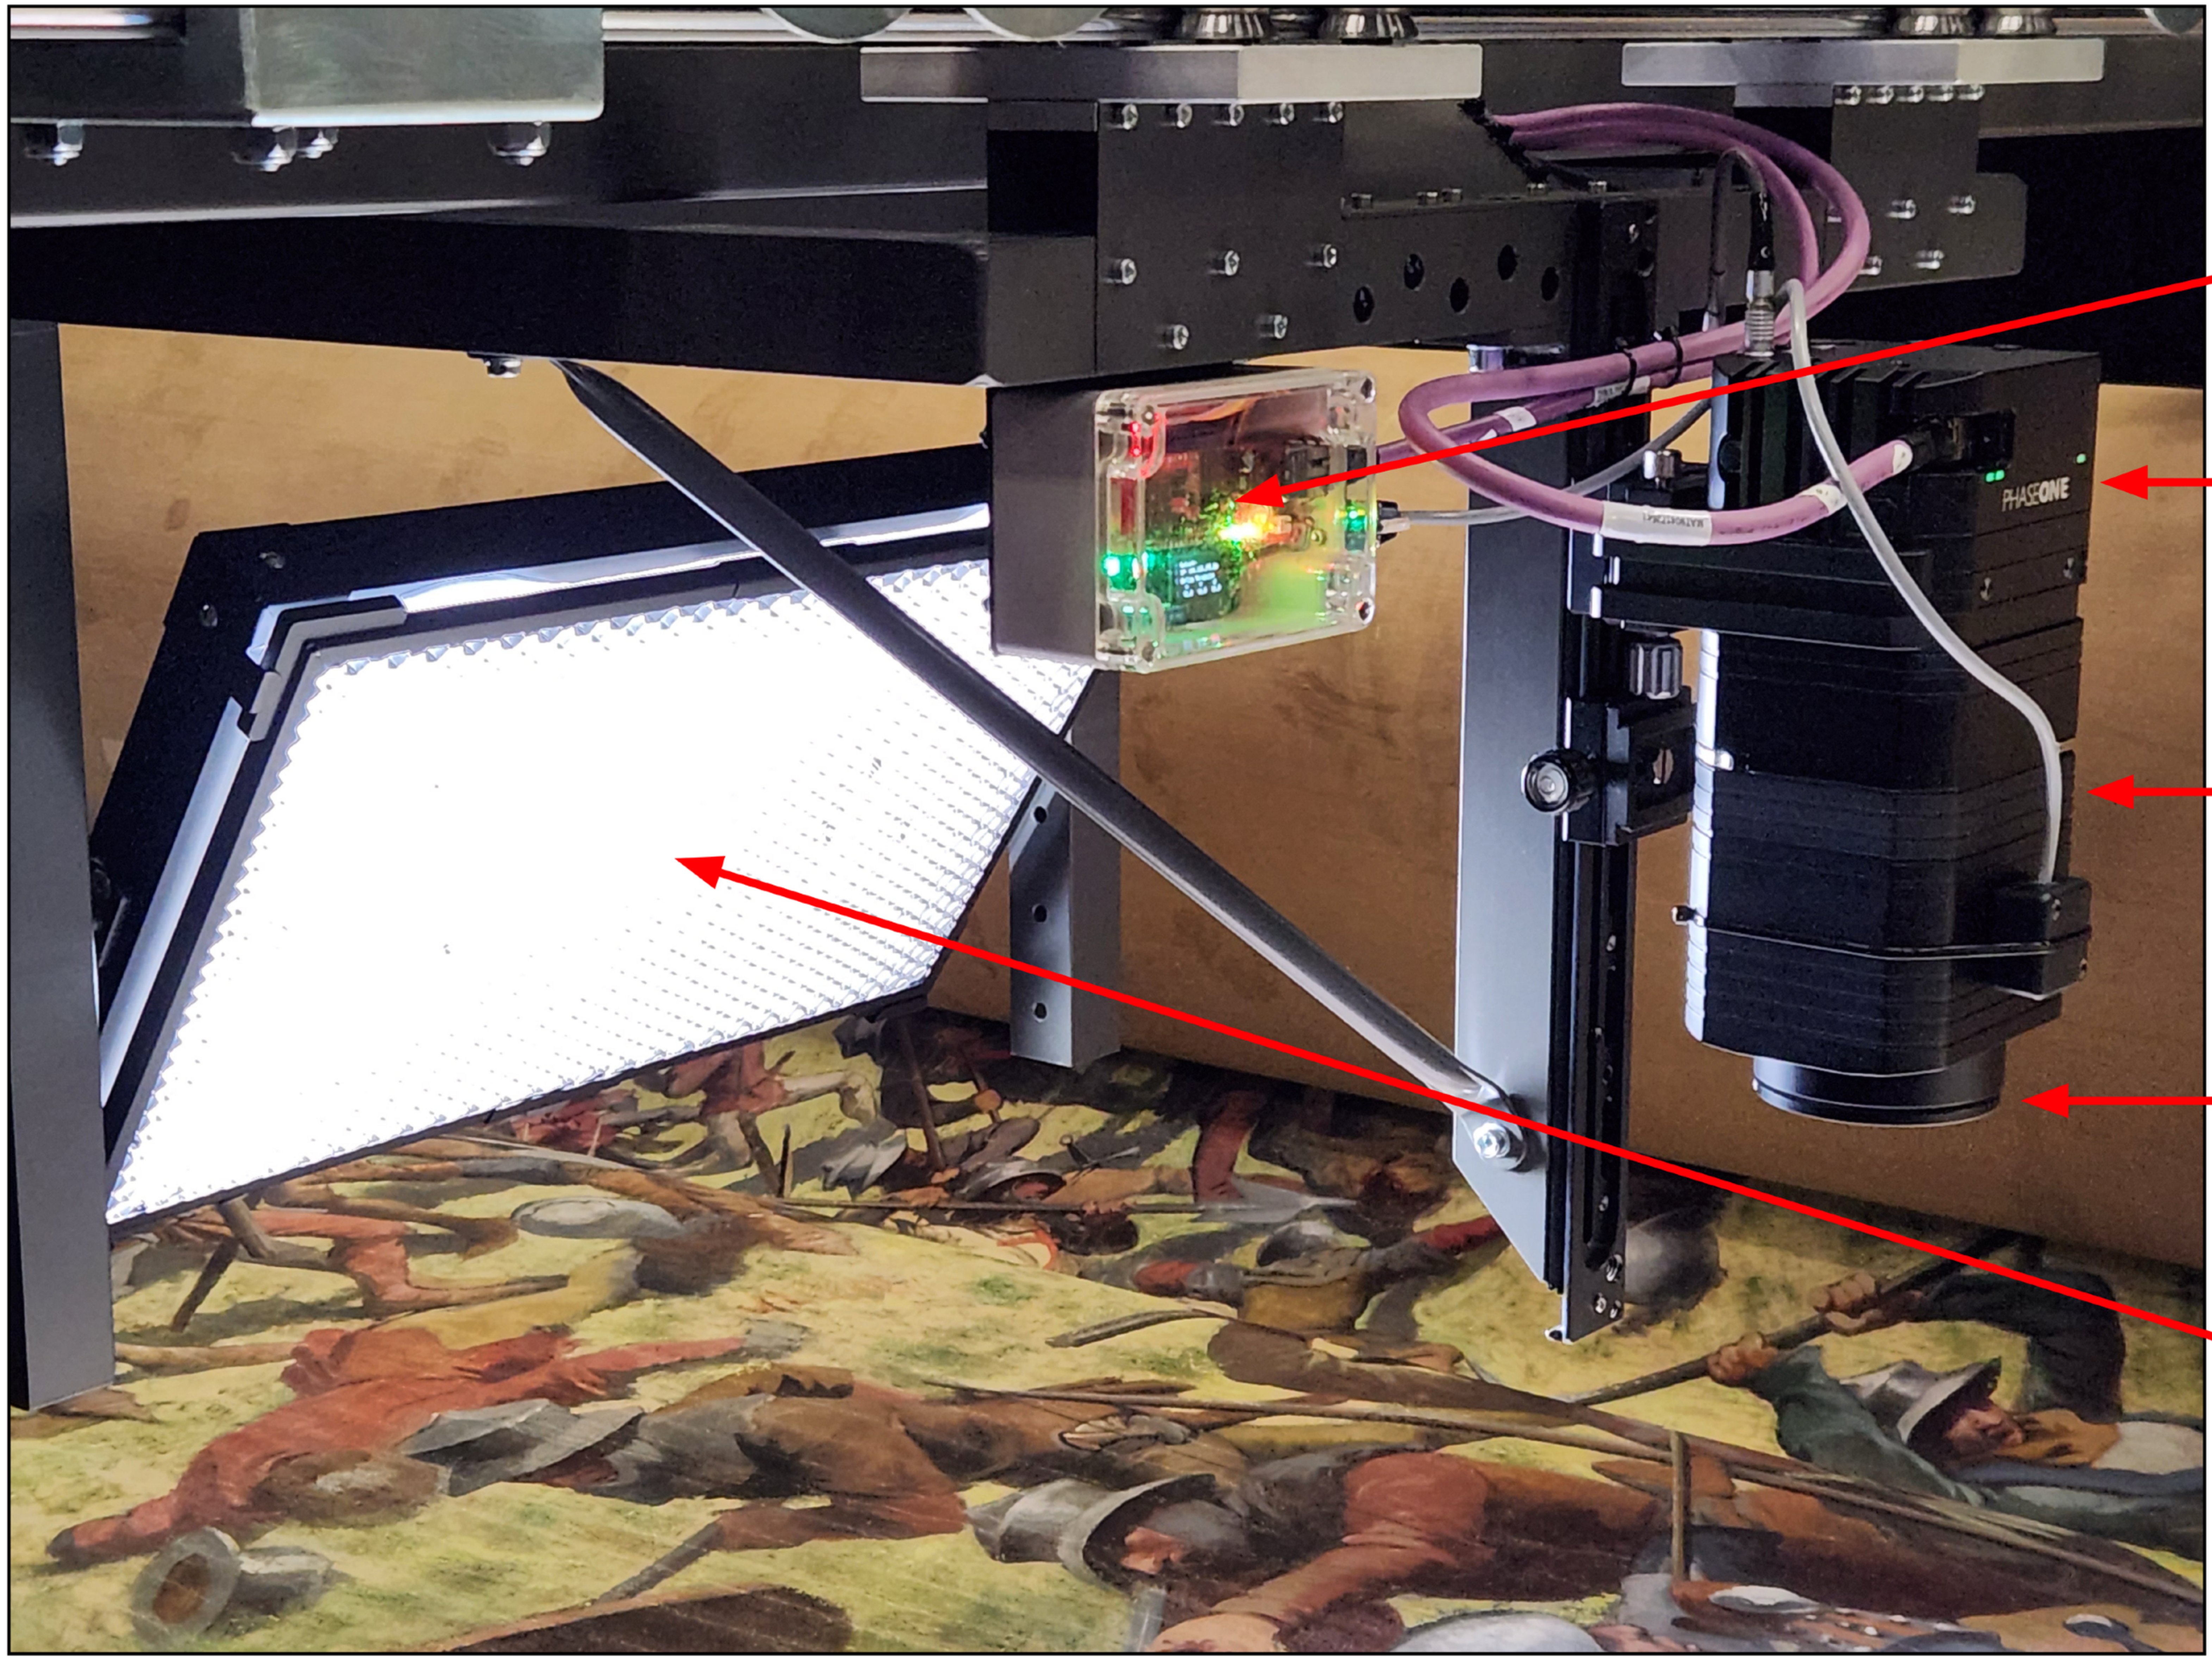

Imaging process

- At the start of each 1m wide strip move to the position of the colour chart, flat field sheet and polarisation knob. Capture and save the three images.
- Move the rig to each grid position, 4 columns x 59 rows
    - Form the photograph name, global x,y addressing
    - Trigger camera initiate transfer/save
    - Calculate next position, zig-zag to minimise travel time
    - Move the camera rig
    - Wait for vibration sensor before proceeding
- Key time saver is moving to next position while file transfer is occurring.
- Perform stitch of current strip with last strip.
- On successful stitch, advance the scroll to the next 1m strip.

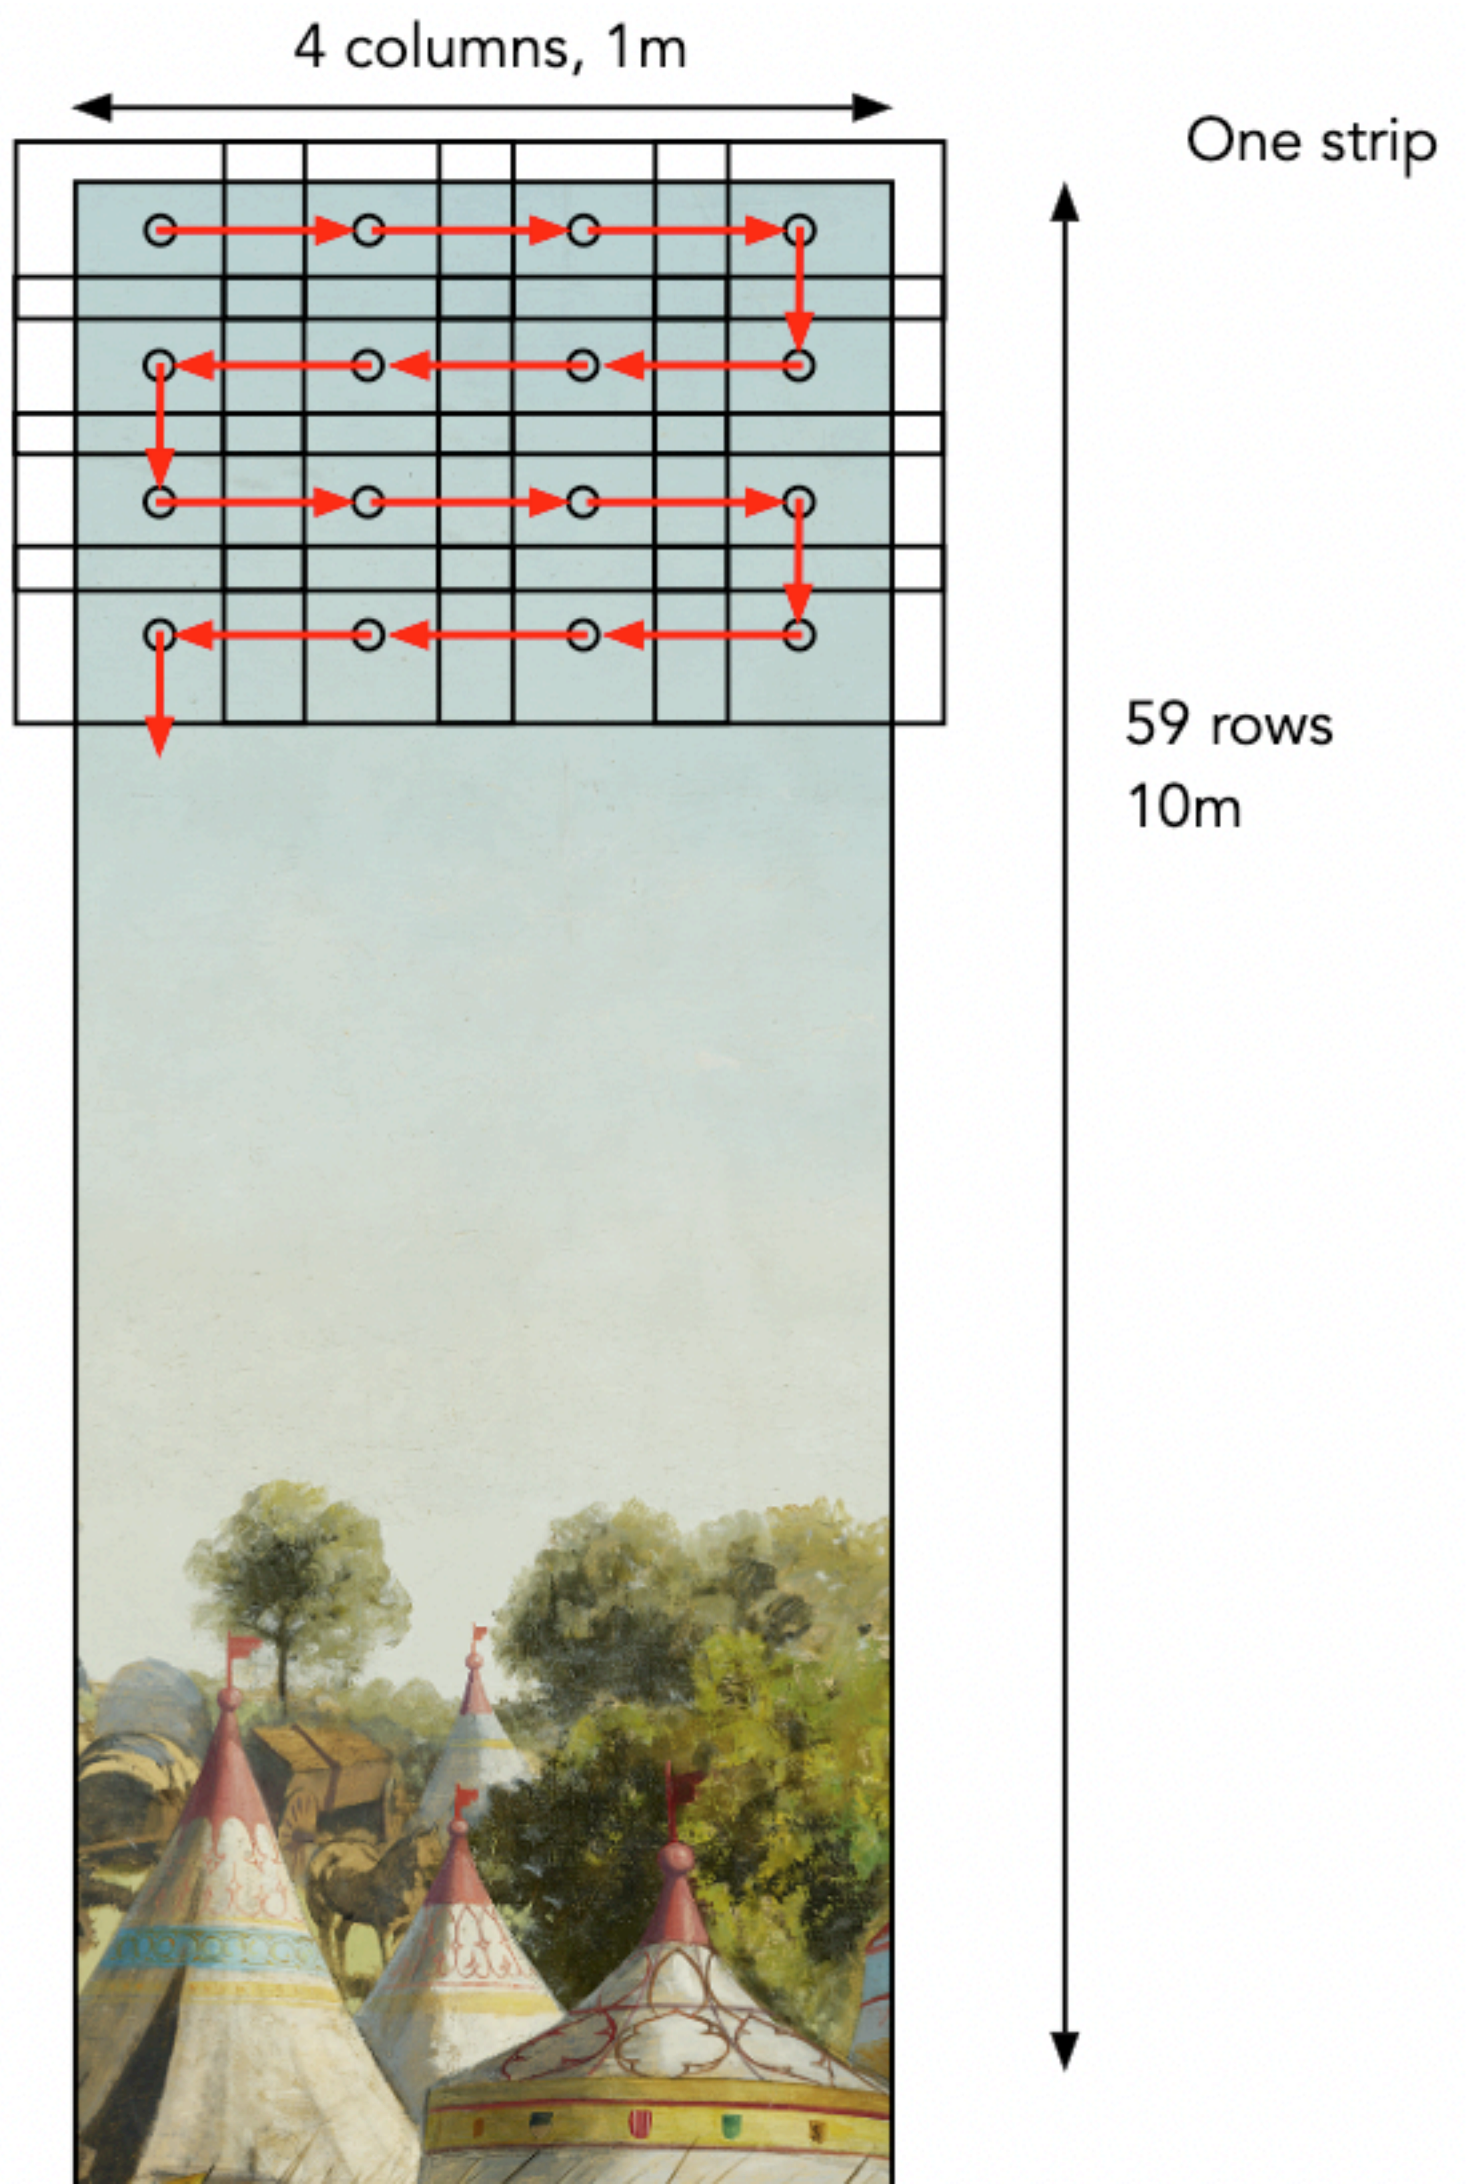

Special images captured once per strip

- Flat field provides correction for the raking light, higher illumination closer to the light
- Colour chart for gamma correction and 1st order colour correction
- Metal orb to check polarisation film isn’t degrading, and lens filter isn’t rotating

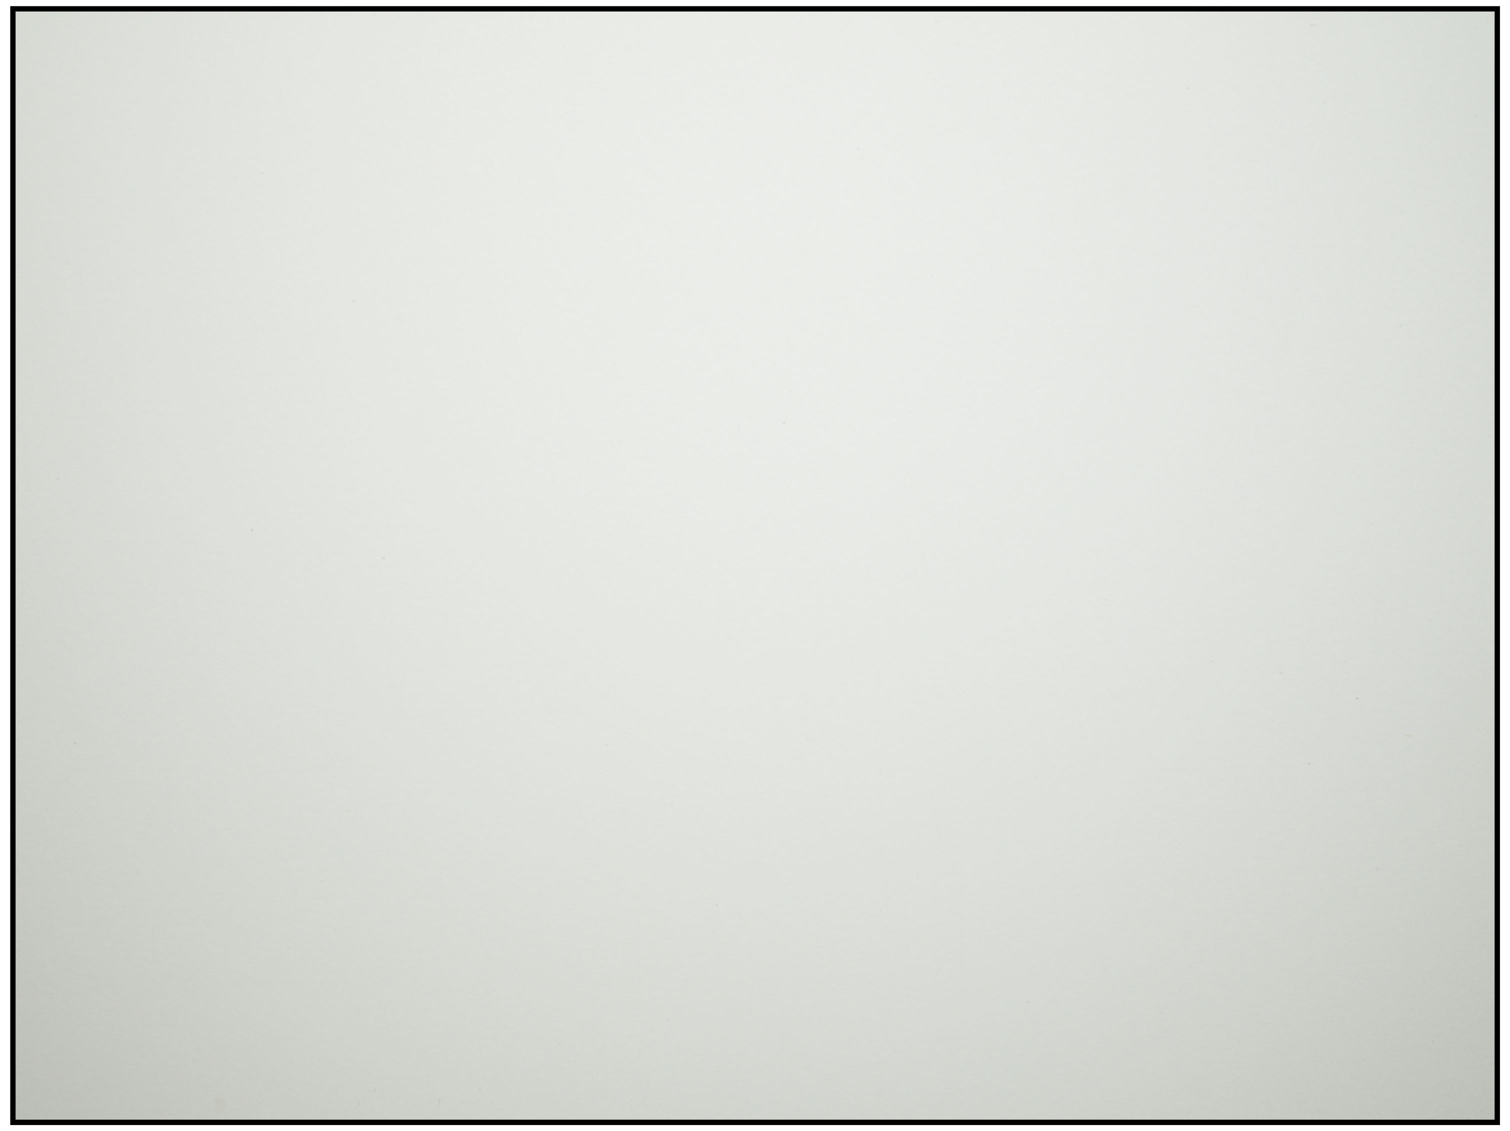

Flat field image

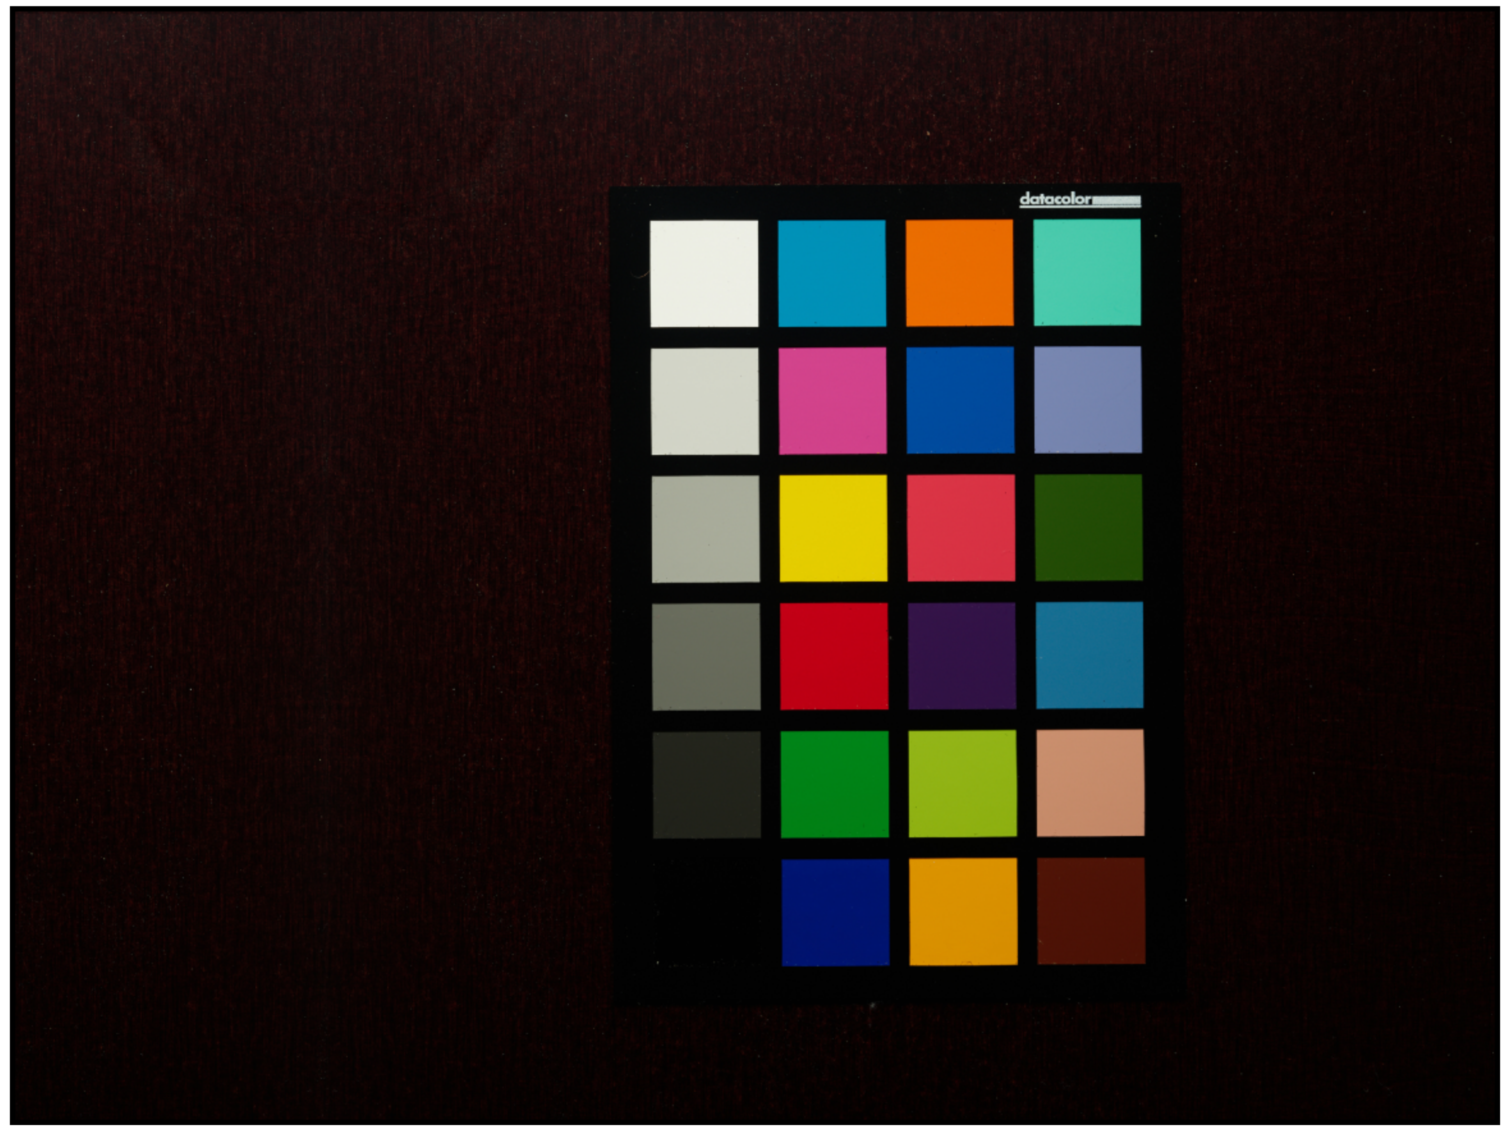

Colour chart

Metal polarisation check
(door handle)

Stitching check

- 50% overlap on image long axis, 30% overlap on image short axis
- Once the scroll is advanced it would not be possible to rewind precisely enough
- Need to ensure no images are missing and the stitching will be successful
- Before advancing the scroll (manually) the current strip is stitched with the last strip
- Approximately 30 minutes to scan a strip, approximately 30 minutes to perform stitch test

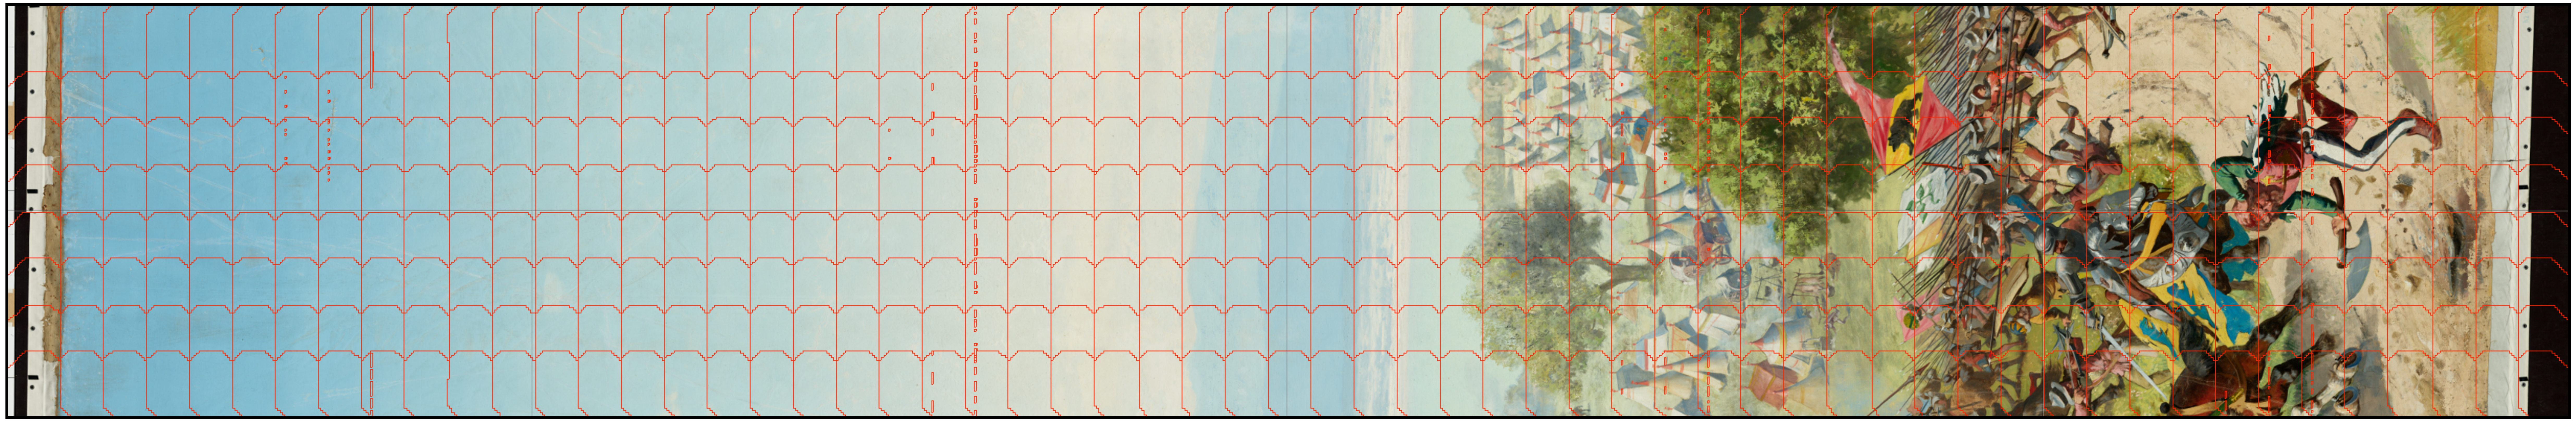

Before advancing the scroll, the current strip (4 columns) are stitched with the previous strip

Pipeline Overview

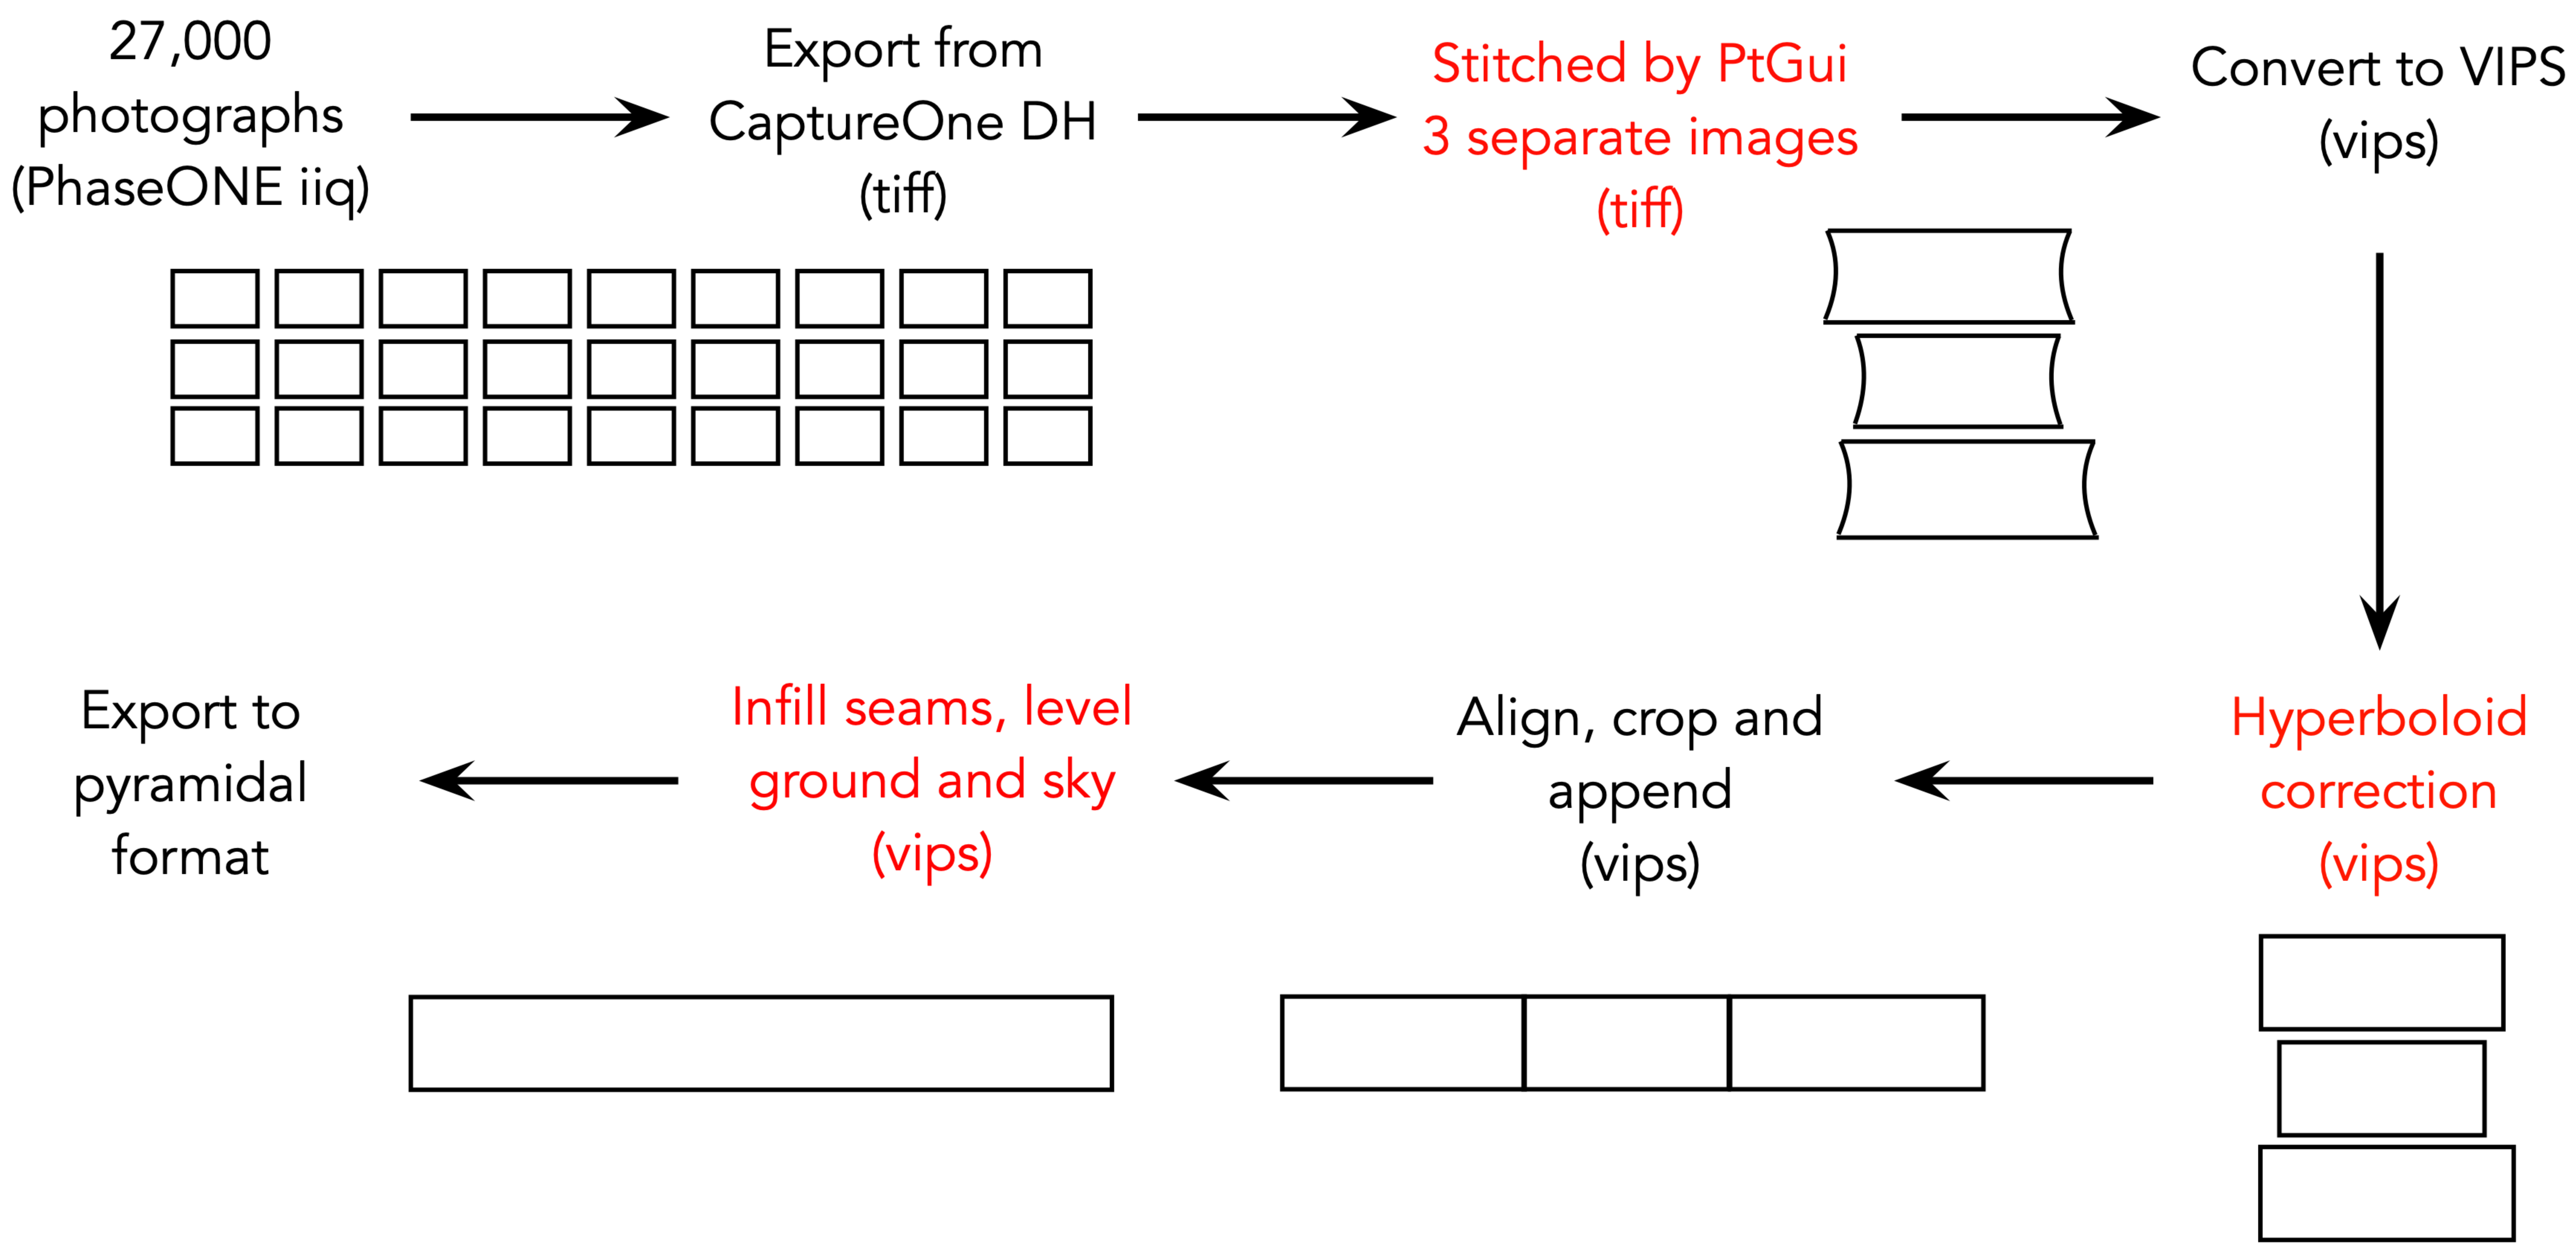

Final imaging statistics

- 27,246 photographs
- Each photograph 14204 x 10652 pixels, 200MB each as iiq raw image (5.45TB in total)
- Acquired over 6 weeks. (7 or 8 working days per scroll, 1 or 2 days between scrolls)
- Final image: 3805340 x 425000 pixels, 8 bit
- Final image file: 1.6 Terapixels uncompressed = 4.8 Terabytes (3 x 8bit RGB)
- Haven't been contradicted yet, "The largest digital image of an artwork"

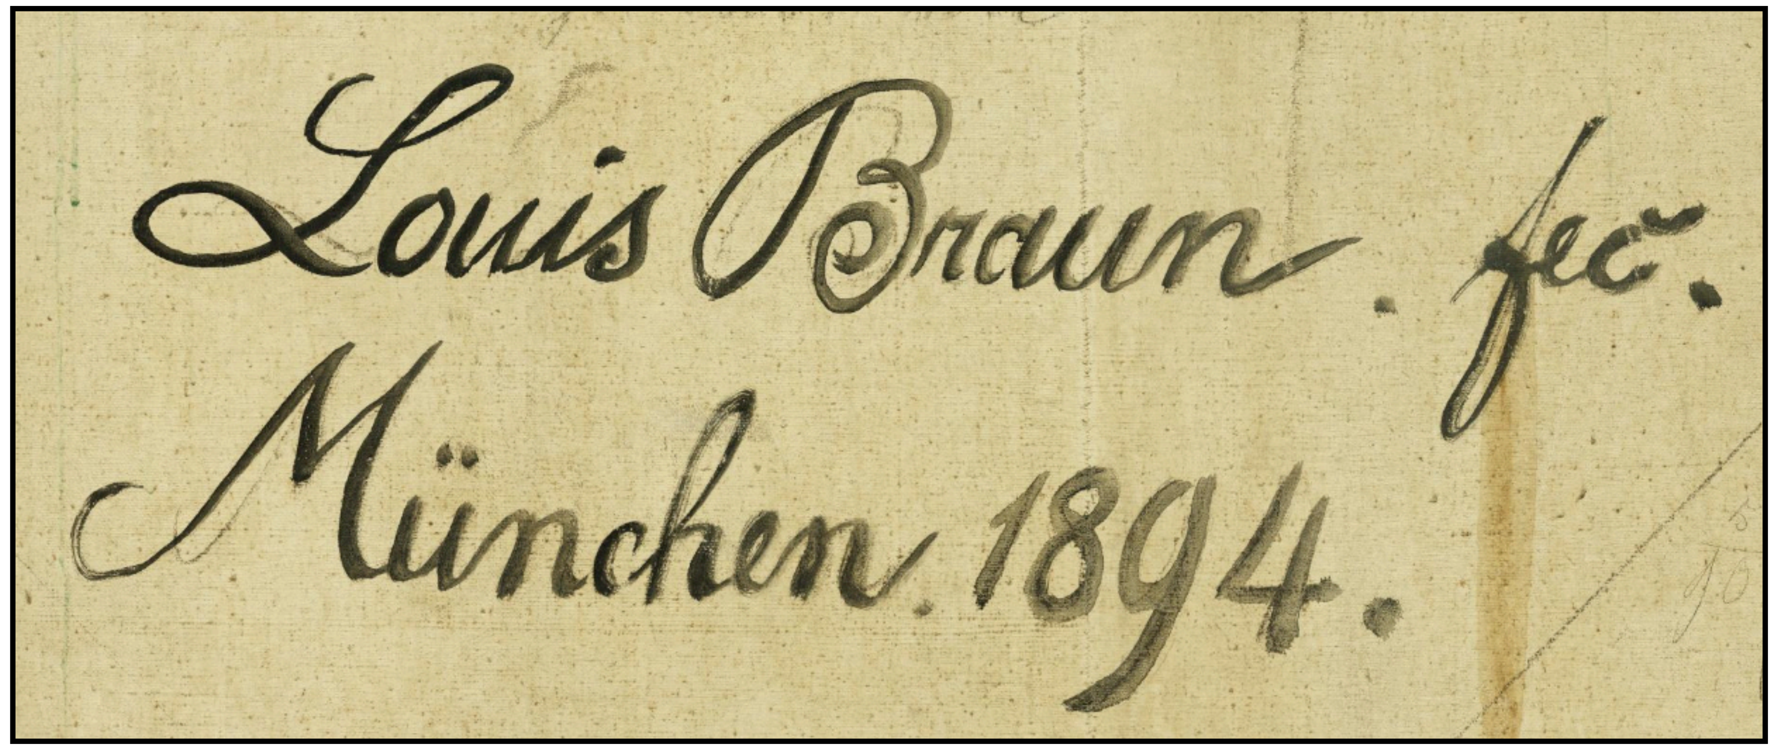

Digitising the Murten Panorama