Tattoo Arithmetic: Measuring the Ink of Prof. Tom ‘TomRocksMaths’ Crawford

Avesta Afshari-Mehr
Department of Computer Science, Durham University*

Figure 1: Both a motivation to study mathematics – and work on abs...

1 Introduction

We’ve all been entranced by the wondrous array of tattoos present across Prof. Tom Crawford’s body: from mathematical equations such as the Euler-Lagrange Equation to Spongebob’s House (with my favourite being the Caterpie!); it’s always been such a satisfying sight for us all. As a keen mathematician and computer scientist, the presence of so many tattoos across our favourite professor’s body does raise a riveting question: just how much of his upper body is covered in tattoos? Could mathematics, or rather, a clever computational approach, tell us the percentage of his body adorned with ink?

Now, you might think of this as some sort of trivial exercise: “just count black pixels and divide by total pixels.” But if only it were so simple! Tom’s tattoos are a beautiful array of colours, including whites and faint grays - while his skin varies subtly with lighting; shadows; and the occasional camera flash. Shadows, in particular, are certain tricksters; mimicking the black ink in edges whilst vanishing under thresholds, creating illusions that the back is covered in far more or far less tattoo than it truly is. And thus began my journey, equal parts detective work and playful maths.

*Also affiliated with the Department of Physics, Durham University.

The aim was to determine a reasonably accurate percentage of tattoo coverage on both his back and front. Yet, even defining what counts as "tattoo" turned out to be more subtle than expected. Should faint shading count? What about white ink that is barely distinguishable from skin? Or thin outlines that disappear under certain lighting conditions? These questions transformed what initially seemed like a straightforward counting exercise into something closer to an investigation, where each decision subtly influenced the final result.

Moreover, the available images introduced their own complications. The back was captured in relatively consistent lighting, while the front involved clothing, shadows, and angles that obscured parts of the body. This meant that a single, uniform approach would almost certainly fail. Instead, the task demanded a more thoughtful strategy: iteratively refining the method, testing assumptions, and visually validating each result to ensure that the algorithm captured tattoos rather than artefacts of lighting or perspective.

2 Methodology

I started with high-resolution images of Tom's back and front. The back was easier: one clean photo sourced from an Instagram post (Tom Rocks Maths, 2025) that showed nearly the entire surface, while the front was acquired from a YouTube short (Tom Rocks Maths, 2022), and that required some clever masking of clothing and underwear to avoid false positives.

For the back, my first approach was the classic computer vision recipe: edge detection. `Canny` edges highlight transitions in intensity, and contours allow us to fill in tattooed regions. The first run was a disaster – it returned a coverage of 90.99%. Almost the *entire* back was flagged as tattoo. Surely, Tom does not yet possess 91% ink coverage; even he and his most dedicated fans would find that implausible (for now... this paper might well provide ample inspiration to achieve such a mission!).

It was here that I had the satisfying revelation: shadows and subtle skin variations were fooling the algorithm. Edges responded to everything, from the gentle shading of his shoulder blade to the faint crease where the arm meets the torso. The solution? Treat different regions differently, like a bespoke tailoring of image processing thresholds.

2.1 Front Challenges

Figure 2: Progression of front tattoo detection from initial segmentation to final refined result.

The front presented new puzzles: shadows from clothing, underwear masking, and tattoos that were faint or light-coloured. Here, a one-size-fits-all thresholding approach failed. My solution was therefore multi-layered and considerably more complex than for that of the back:

- **Contrast enhancement:** Contrast Limited Adaptive Histogram Equalisation (CLAHE) improved visibility of faint tattoos.
- **Region-specific edge detection:** Different thresholds for "problem" and "high-sensitivity" regions.
- **Exclusions:** Masking underwear and other irrelevant areas prevented misclassification.
- **Manual adjustments:** Very light tattoos and a few tricky rectangles were marked manually to ensure complete coverage.

The front ultimately became a balancing act between automation and precision. Too sensitive, and shadows appeared as tattoos; too conservative, and faint tattoos disappeared altogether. Carefully tuning each region eventually led to a stable and visually convincing segmentation.

2.2 Back Challenges

Although the back initially seemed simpler, it introduced its own unique challenges. The most significant issue arose from lighting gradients across the skin. Areas near the shoulders appeared darker due to shadowing, while lower regions had more uniform illumination. These differences caused edge detection to behave inconsistently across the image.

Two areas of the back caused the most trouble: the bottom-right, where skin was mistaken for tattoos, and the bottom-left, where tiny black-ish tattoos hid in plain sight.

Bottom-right region: The initial algorithm over-detected tattoos. By increasing the thresholds for edge detection and filtering out small contours, I could dramatically reduce false positives. Essentially, I made the algorithm more conservative in this zone, letting only the true edges survive.

Bottom-left region: Here, the opposite problem occurred: actual tattoo ink was being ignored because it was dark but low-contrast against the skin. The solution was cleverer: adaptive thresholding *plus* morphological operations. Adaptive thresholding allowed my algorithm to detect subtle intensity differences, while morphological opening and closing removed noise without erasing the tiny tattoos. This careful, region-specific processing ensured the back mask was finally accurate.

Another complication came from tattoo density. Some regions featured large, bold tattoos, while others contained delicate line work. A global threshold struggled to accommodate both simultaneously. Increasing sensitivity helped capture delicate tattoos but also introduced false positives along muscle contours and natural skin shading. Conversely, reducing sensitivity eliminated false positives but caused finer tattoos to vanish.

To address this, my algorithm evolved into more of a hybrid approach. Large-scale edge detection captured the dominant tattoos, while localised refinements targeted areas where the global method failed. This iterative refinement process gradually reduced errors and improved overall accuracy.

I soon realised that tattoo coverage is not *just* about detecting edges or ink; it's about understanding context. A shadow isn't a tattoo. A patch of clothing isn't a tattoo. And yet, the algorithm must learn to make these distinctions without being told explicitly.

3 Results

The development of my results tells a fascinating story.

3.1 Back Results

The final segmentation for the back is shown below. The overlay highlights detected tattoo regions in red, while the mask displays the binary segmentation used for pixel counting.

Figure 3: Back segmentation: overlay (left) and binary mask (right).

Further, the back initially produced highly exaggerated values:

- 90.99% – Initial naive segmentation;
- 41.38% – First refinement;
- 31.83% – Bottom-right correction; and,
- 38.30% – Final refined result.

The jump from 31.83% to 38.30% reflects the successful detection of small tattoos previously missed in the bottom-left region.

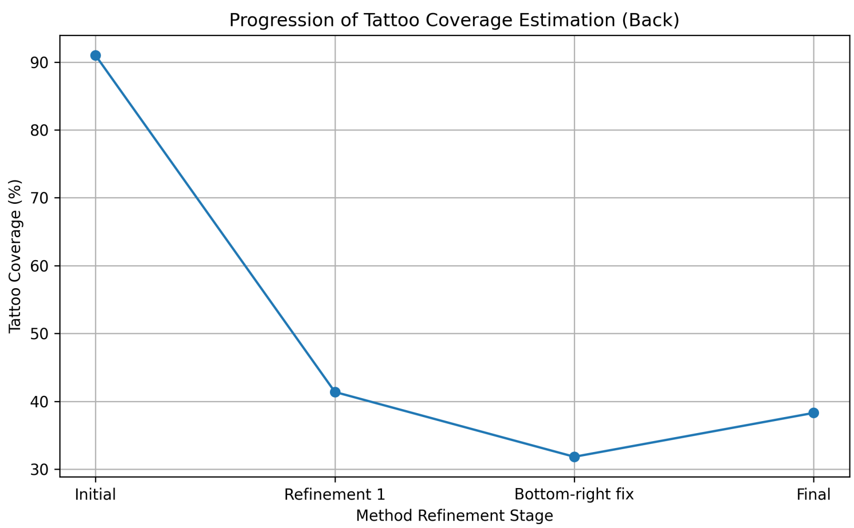

Figure 4: Progression of tattoo coverage estimation for the back. Each refinement stage improved segmentation accuracy and reduced false positives.

3.2 Front Results

The front segmentation required additional refinements:

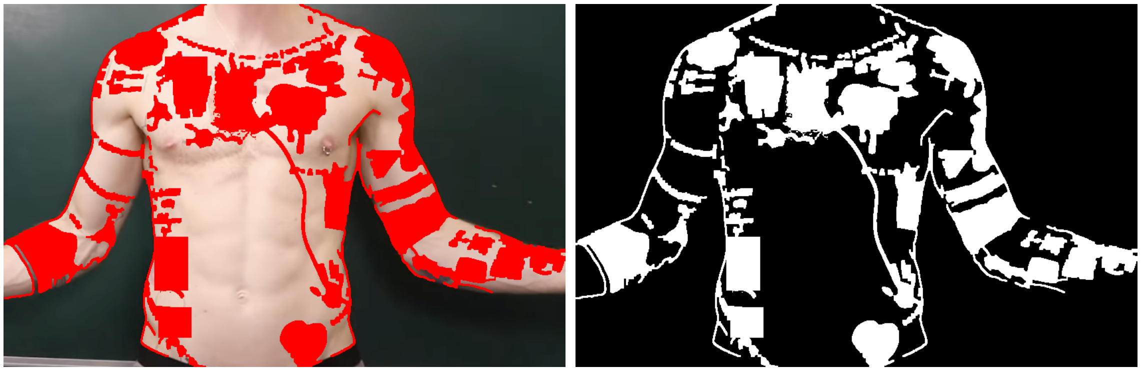

Figure 5: Front segmentation: overlay (left) and binary mask (right).

Unlike the back, the front showed more gradual refinement:

- 22.03% – Initial estimate;
- 22.77%;
- 22.83%; and,
- 22.91% – Final result.

Unlike the back, the front required fewer dramatic adjustments barring the forced covers (shown as rectangles), which reflects the relatively stable segmentation once exclusions were properly implemented.

3.3 Combined Coverage

To estimate total upper body tattoo coverage, we combine the front and back values. For simplicity, the front and back were assumed to represent approximately equal surface areas, and the overall coverage was estimated using their average:

$$\text{Average coverage} = \frac{38.30 + 22.91}{2} \tag{1}$$

$$= 30.61\% \tag{2}$$

Thus, approximately **30.61%** of Prof. Tom Crawford's upper body is covered in tattoos. Nearly *one-third* of Tom's upper body is adorned with tattoos!

4 Discussion

While the results are compelling, several sources of uncertainty remain. One of the most significant limitations arose from manual adjustments. In some cases, particularly with faint tattoos, rectangles were manually applied to ensure inclusion. While effective, this approach introduces uncertainty, as these rectangles inevitably include small areas of skin. This artificially increases the measured tattoo coverage.

Similarly, some regions were deliberately excluded to avoid false positives. While necessary, this creates the opposite effect: genuine tattoos near these areas may have been omitted. These competing sources of error introduce an unavoidable degree of uncertainty.

Other sources of uncertainty included:

- Skin-to-background edges being incorrectly flagged as tattoos;
- Armpit shadows appearing as dark ink;
- Edges between skin and underwear; and,
- The microphone and cable visible in the image.

Each of these contributed small inaccuracies to the segmentation. Estimating their combined impact suggests an approximate uncertainty of $\pm 2.77\%$ in the final result.

Another interesting observation emerged during refinement: tattoos with bold outlines were consistently easier to detect than shaded tattoos. This highlights an unexpected intersection between artistic style and computational detection. Bold tattoos are mathematically easier to measure.

Future improvements could include machine learning segmentation models, colour clustering approaches, or multiple images under consistent lighting conditions. These would reduce manual intervention and further improve accuracy.

5 Conclusion

What began as a playful question evolved into a surprisingly rich mathematical investigation. Through computer vision; iterative refinement; and pixel-level analysis, we estimated that approximately 30.61% of Prof. Tom Crawford's upper body is covered in tattoos, with an uncertainty of roughly $\pm 2.77\%$.

Beyond the mathematics itself, the journey we've ventured reveals something much deeper: mathematics (and its subsequent computations) can quantify even the most unexpected things. From Caterpie to Euler-Lagrange, Tom's tattoos became not just art, but rather data: patterns just waiting to be measured.

And in the end, each percentage isn't just a number – it's a tiny victory over shadows; thresholds; and stubbornly faint tattoos that could not be identified algorithmically, despite my greatest efforts. What started as a cheeky question about ink turned into a surprisingly satisfying mathematical adventure.

6 References

- Tom Rocks Maths (2025). "The 3rd most popular on Tom Rocks Maths in 2025 is... Tattoos." Available at `https://www.instagram.com/p/DS0Zpn0jYmu/?img_index=3` (Accessed: 6 April 2026).

- Tom Rocks Maths (2022). "Meet the Naked Mathematician." Available at `https://youtube.com/shorts/uCQWcICuDa0?si=iLkqgDhTdmxw4y7p` (Accessed: 6 April 2026)