capstract:2611.00046v1 [cs.CV] 20 Nov 2026

Establishing quorum from coffee-ring impressions

Mara Velin
Visual Measurement Group
Hather Quay Institute

Abstract

Five coffee rings remain after a committee meeting. To determine whether five members were present, a secretary must first distinguish five mugs from four mugs and a refill. We study this problem using synthetic contact prints from 40 eight-member committees, each with a gallery of enrolled mugs. Matching the angular variation of the contact radius identifies the source of 96.1% of complete prints and 75.3% of half prints. We then vary who carried the mugs while keeping the images fixed. With four members present, recovering the full circumference instead of one quarter reduces false quorum certificates from ten to one in 40 meetings with a refill. The same improvement raises false certificates from nine to 35 when a member brings an absent colleague's mug. A source oracle certifies all 40 meetings in the latter condition. We give the paired seating arrangements that make human quorum unidentifiable even when the tableware register is complete.

1 Introduction

A committee with eight members needs five to conduct business. After the room has been cleared, five rings remain on the table. The secretary would like to know whether to enter a quorum in the minutes. Counting rings gives a poor answer. One member can put a mug down twice. A more promising procedure is to keep a specimen impression from each member's mug and strike out duplicates by visual comparison. This procedure requires an enrollment session but dispenses with the attendance sheet at all subsequent meetings.

We examine what the secretary could recover under favorable imaging conditions. Each member receives one mug for the duration of the study. The mugs are similar enough in diameter that size alone seldom settles a match. Small departures of the foot from a circle persist from one impression to the next. We render these departures into grayscale images and ask a matcher to identify each source from the eight specimens on file. The matcher sees the same pixels that would be available to the secretary, including a blank sector where the table has been wiped.

The question of record comes after retrieval. An impression has a source mug, a registered owner and a person who brought it into the room. These labels usually coincide in an orderly service. We change the last one while holding the first two fixed. This gives an attendance test in which the evidence submitted by a member can be examined in that member's absence.

2 Contact prints as enrollment records

Daugman's iris matcher uses angular registration and masks out sectors unsuitable for comparison [1]. A mug foot offers a simpler annulus. At each angle we retain one number, the radius of deposited material. Its variation across angle carries the identifying information. A shift in angle costs the member nothing, since no rule specifies which way the handle must face during proceedings.

Instance retrieval benchmarks distinguish gallery objects from the views used to query them and specify which matches qualify as relevant [3]. Here relevance has two possible definitions. An im-

Preprint. Under review.

pression made by the treasurer’s mug is a positive query for that mug on every occasion. It is a positive query for the treasurer only on occasions the treasurer attends. We score both tasks, using the same gallery for each.

Our prints model direct transfer from a wet ceramic foot. An evaporating drop can also leave a ring, through transport to its pinned contact line [2]. That process gives no reason for the perimeter of a spilled drop to match a mug in the cupboard. We admit only foot impressions and provide the crop around each one. Finding such crops among drops, spills and saucer marks would be a separate detection task.

3 Image formation and retrieval

We drew 40 independent cabinets, each containing eight mug profiles. In a 96×96 pixel crop, the reference radius of mug i at angle θ is

$$r_i(\theta) = R_i + \frac{0.85}{\sqrt{10}} \sum_{k=3}^{12} (a_{ik} \cos k\theta + b_{ik} \sin k\theta).$$

Here R_i is normal with mean 33 pixels and standard deviation 0.9. All Fourier coefficients are independent standard normals. The amplitude in this expression sets the between-mug variance. On each contact we add a radial offset of standard deviation 0.35 pixels and an independent angular field with harmonics 2 through 13 and pointwise standard deviation 0.38 pixels. The latter term allows water to shift the line even when the mug has been placed exactly as before. These are specified image conditions. We have no estimate of their distribution in manufactured tableware.

For each impression we drew a translation of up to 1.5 pixels per axis and a uniform rotation in steps of one angular sample (1/128 turn). The line has a Gaussian cross section of width 1.05 pixels (standard deviation). At each bearing we set the peak darkness to 0.52 plus a fresh angular field with harmonics 1 through 6 and standard deviation 0.11, clipped to the interval [0.22, 0.82]. The background has mean 0.89 and a horizontal slope drawn uniformly between −0.04 and 0.04 gray units per half image width. We added independent normal pixel noise of standard deviation 0.027 on a gray scale from 0 to 1, then quantized to eight bits. Each cabinet supplied eight full enrollment prints and two fresh query impressions per mug. We retained four nested sectors of every query, with a common random starting angle. Thus each coverage condition has 640 queries. The 320 enrollment images are separate draws. Figure 1 shows the first query in cabinet order.

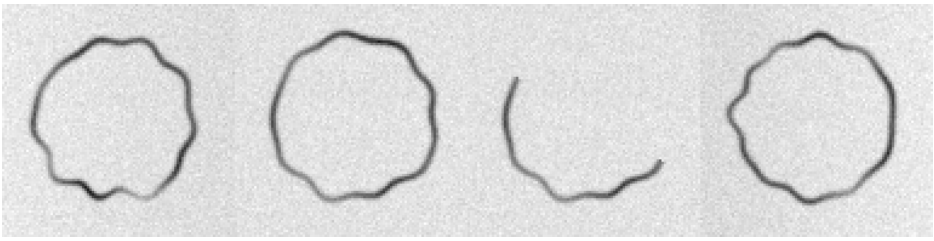


Figure 1: From left, the enrollment impression for mug 1, a new full impression from mug 1, its half-arc query, and the enrollment impression for mug 2. All four crops have the same pixel scale. The two middle panels share a contact event.

The wipe acts on the deposited material before the camera noise is added. The erased sector therefore contains the same background fluctuations as the interior and exterior of the ring. A method can use the two ends of a partial arc to estimate the extent of the wipe. Their location is drawn anew for each contact and bears no fixed relation to the mug handle.

3.1 Reading a partial arc

At 128 bearings from the crop center we sampled radii from 25 to 40.75 pixels at quarter-pixel intervals by bilinear interpolation. We located the darkest sample, took the darkness-weighted radial centroid within two pixels of it, and retained the bearing if its peak darkness relative to a gray level

of 0.85 exceeded 0.16. The extractor receives neither the rotation nor the wipe mask. A fit of a constant and first angular harmonic removes mean size and, to first order, center displacement.

Write y for the query radius trace, g_i for an enrolled trace and V for the visible query bearings. Let P_V take residuals after least-squares regression on $1, \cos\theta, \sin\theta$ at those bearings. Our score is

$$d(i) = \min_{s=0,\dots,127} \frac{1}{|V|} \left\| P_V \left(y_V - (g_i)_{V+s} \right) \right\|_2^2 .$$

Indices wrap around the ring. We assign the mug with the smallest score. Fitting both traces on the observed sector matters because a single short arc admits several plausible centers. Fitting the gallery on a whole circle and the query on a quarter circle would give the two residuals different reference circles.

We compared three alternatives. The radius rule uses the closest fitted intercept. A rotation-invariant descriptor takes the magnitudes of Fourier coefficients 3 through 12 of each residual, filling missing bearings with zero and dividing by the observed sample count. It selects the nearest descriptor in Euclidean distance. The fixed-angle ablation uses the registered score with $s = 0$. All four rules use the same extracted traces. Chance accuracy in a cabinet is 12.5%.

4 Results

4.1 Mug identification

Full rings allow the matcher to recover the correct mug on 615 of 640 impressions. Half a ring suffices on 482. Figure 2 gives the coverage dependence. At full coverage, keeping Fourier magnitudes alone costs 17.2 percentage points. A magnitude spectrum cannot distinguish two profiles with identical harmonic amplitudes and different relative phases. At quarter coverage it has the additional problem that the two edges of the retained sector appear in the spectrum. The registered matcher still wins there, though only 170 queries out of 640 return to their owner.

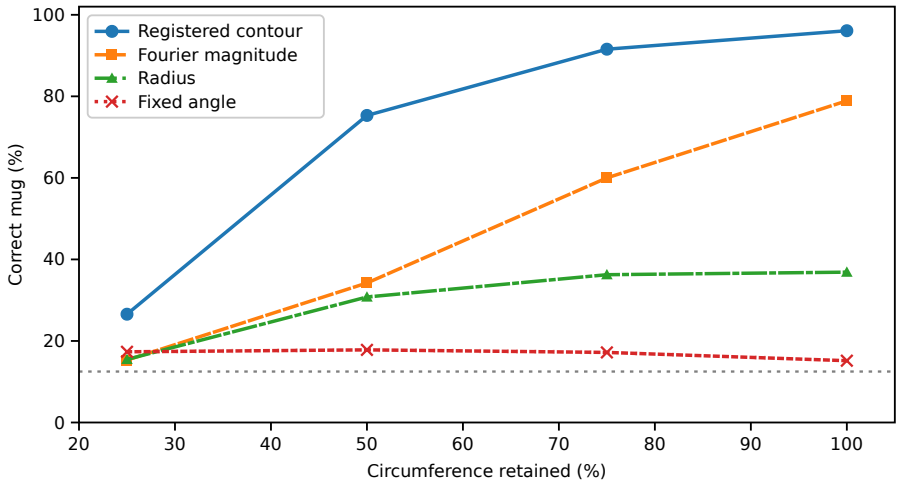


Figure 2: Top-1 mug identification in eight-mug cabinets. Each point comprises 640 query views. Coverage conditions reuse the same contacts. The dotted horizontal line gives uniform guessing.

Across 5,000 bootstrap resamples of whole cabinets, the registered matcher's 95% percentile intervals were 94.1 to 97.8% for full rings, 71.9 to 78.6% for half rings, and 23.4 to 29.7% for quarter rings. The cabinet is the sampling unit because every query in it faces the same seven alternatives. Size alone recovered 36.9% on complete prints, above the 15.2% from comparing contours at the angle in which each mug happened to land. On short arcs, part of the distinguishing shape can be absorbed into the fitted center. The secretary has the least evidence of identity precisely where the center estimate is least constrained.

4.2 Circular-foot control

A cabinet of unusually round mugs presents a harder test. We rendered the same contacts twice more, halving and then removing the stable angular deviation of each foot. The mean radii, transient distortion fields and image noise draws were held fixed. Table 1 gives the full-ring results. With circular feet, shape matching falls within sampling error of chance even though size still identifies over a third of the mugs. In this cupboard a caliper reading would be preferable to a contour record.

Table 1: Full-ring identification (%) as a function of persistent radial variation. Each column uses 640 prints. The bracketed values are cabinet bootstrap intervals for the registered rule.

Retrieval rule	0 pixels	0.425 pixels	0.85 pixels
Radius	36.9	37.2	36.9
Fourier magnitude	12.2	34.1	78.9
Registered contour	13.8 [11.1, 16.3]	46.7 [43.4, 50.2]	96.1 [94.1, 97.8]

4.3 Taking the roll

Within each cabinet we drew five different mugs without replacement. Each left one query print. In the ordinary meeting these mugs belonged to the five members in the room. For the carried-mug condition, four of those members attended and one carried in the fifth member’s mug. The registration ledger continued to assign that mug to the absent member. We reused the same five images in the same order so that illumination, placement and contact quality were paired exactly.

A second table setting contained only the first four mugs. The first mug made an additional, independently rendered impression, giving the secretary five marks to inspect as before. Four members taking a refill and five members sharing four mugs could each leave that record. For all four attendance conditions the rule certified quorum if the five prints returned five distinct gallery identities. It made no use of the order of refreshments.

Table 2 reports how many of the 40 boards certified quorum. The raw count rule certified all 40 in every column. Identification removed 39 of those certificates when four mugs left five complete rings. It retained 35 when four people supplied five mugs. In this latter condition, the five failures to certify arose when different mugs received the same name. Each such collision corrected the minutes.

Table 2: Meetings certified as having five members present, out of 40. The two columns within each mug count use identical images. The oracle reads the true mug identity of every print.

Evidence read	Five attend	Four, extra mug	Four, refill	Five, shared mug
Marks counted	40	40	40	40
Quarter arcs	9	9	10	10
Half arcs	21	21	6	6
Three-quarter arcs	35	35	1	1
Full rings	35	35	1	1
Mug oracle	40	40	0	0

Improved recognition has opposite effects on the two kinds of inquorate meeting. Extending quarter arcs to full circles withdrew nine erroneous certificates from refill meetings and issued 26 additional ones to meetings supplied with an absent member’s mug. Perfect mug recognition would give the absent member credit in the remaining five.

5 What the minutes identify

Let X denote the five images and their enrollment gallery. Within either pair of seating arrangements, every realization of X occurs once with a quorum and once without. Any decision based on X must return the same answer for both. With equal weight on the two arrangements its attendance error is exactly one half, including for a randomized rule in expectation. This is a property of the paired experiment. It does not estimate the frequency with which committee members bring someone else’s crockery.

There is a second ambiguity in the performance measure. A clerk could move every name in the enrollment ledger down one row, wrapping the last to the top. Every correctly identified mug would then be reported under a different member’s name. Every entry in Table 2 would survive the amendment because the number of distinct names is invariant to a permutation. For minutes that list who was present, an audit of quorum decisions alone would miss the entire clerical error.

The synthetic feet have persistent irregularities by construction. A physical test would need repeated wet impressions from different mugs of the same production lot, on different days, to estimate the stable and transient components separately. Our enrollment prints are complete, all contacts appear in isolated crops, and every mug belongs to the cabinet. Overlapping rings and visitors’ mugs fall outside this experiment. Even under the stated conditions, a single tabletop image has no time coordinate. Five members arriving one at a time can leave the same evidence as five members sitting together.

The quantity the secretary can estimate from these images is the number of distinct registered mugs that touched the table. For boards whose rules require the simultaneous presence of five members, the attendance sheet has one field the cabinet cannot supply.

5.1 Computational details

We generated the cabinets using NumPy 1.26.4 with PCG64 seeds 41800 through 41839, one seed per cabinet. A separate seed of 771 governs the cabinet bootstrap. No matcher uses fitted training parameters. Image generation, all four retrieval rules, the meeting decisions and resampling took about four seconds per cabinet series on one CPU thread. Across the three specified radial amplitudes and four coverage settings, the experiment scored 7,680 query crops.

References

[1] John Daugman. How iris recognition works. *IEEE Transactions on Circuits and Systems for Video Technology*, 14(1):21–30, 2004.

[2] Robert D. Deegan, Olgica Bakajin, Todd F. Dupont, Greg Huber, Sidney R. Nagel, and Thomas A. Witten. Capillary flow as the cause of ring stains from dried liquid drops. *Nature*, 389:827–829, 1997.

[3] Filip Radenović, Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondřej Chum. Revisiting Oxford and Paris: Large-scale image retrieval benchmarking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5706–5715, 2018.