w07a: RANSAC, solution

Notation as in the Week 7 slides: e is the outlier fraction, s the number of points per sample, N the number of iterations, d the inlier threshold and p the desired probability of success.

Q1. Iteration count.

A point is an inlier with probability $1-e$. Picking the s points independently, a sample is good (all inliers) with probability $(1-e)^s$ and bad with probability $1-(1-e)^s$. For N independent samples, all are bad with probability $\left(1-(1-e)^s\right)^N$, so at least one is good with probability $1-\left(1-(1-e)^s\right)^N$. We want this to be at least p:

$$\begin{aligned} 1-\left(1-(1-e)^s\right)^N &\geq p \\ \left(1-(1-e)^s\right)^N &\leq 1-p \\ N\log\left(1-(1-e)^s\right) &\leq \log(1-p) \\ N &\geq \frac{\log(1-p)}{\log\left(1-(1-e)^s\right)} \end{aligned}$$

The last step flips the inequality because $0 < 1-(1-e)^s < 1$, so its log is negative. Take the smallest whole number:

$$\boxed{N = \left\lceil \frac{\log(1-p)}{\log\left(1-(1-e)^s\right)} \right\rceil}$$

The number of data points n does not appear.

Q2. The table ($p = 0.99$):

e	0.1	0.3	0.5	0.7
line, $s = 2$	3	7	17	49
homography, $s = 4$	5	17	72	567

For example, $e = 0.5$, $s = 2$: $N = \lceil \log 0.01 / \log 0.75 \rceil = \lceil 16.01 \rceil = 17$.

What changes between the rows. N grows *very* steeply with the sample size, because $(1-e)^s$ shrinks geometrically. At $e = 0.7$ a line needs 49 iterations and a homography 567: about 12 times more for two extra points. This is why RANSAC is cheap for lines and expensive for models needing many correspondences, and why minimal samples matter.

Values within one or two of the table are acceptable if the student did not round up; what matters is the trend and using the formula consistently.

Q3. The inlier threshold d.

- **Far too small:** almost no point is an inlier. Every model scores near zero, so the winner is close to random, and the refit uses only a few points.
- **Far too large:** outliers count as inliers too. Every model scores about the same, and the refit is plain least squares on bad data, the problem RANSAC was meant to avoid.

Good answers relate d to the noise: if the inlier noise has standard deviation σ, set d between 2σ and 3σ. A stronger answer says where σ comes from: the sensor, or an estimate from the residuals near the model using the MAD, $\hat{\sigma} = 1.4826 \cdot \text{median}_i\, |r_i|$.

Q4. Adaptive stopping.

The best model so far explains $n_{\text{in}} = 63$ of $n = 100$ points, so

$$\hat{e} = 1 - \frac{n_{\text{in}}}{n} = 1 - \frac{63}{100} = 0.37.$$

Then, with $s = 2$ and $p = 0.99$,

$$N = \left\lceil \frac{\log(1 - 0.99)}{\log\left(1 - (1 - 0.37)^2\right)} \right\rceil = \left\lceil \frac{\log 0.01}{\log 0.6031} \right\rceil = \lceil 9.11 \rceil = 10.$$

This is the third update in the slides' worked run, which stops at 10 iterations.

Why it is safe. Until RANSAC finds the true model, the best model so far misses some inliers. So n_{in} is usually too small, $\hat{e}$ too large, and N too large: RANSAC does extra work rather than stopping too early. Each better model lowers $\hat{e}$ and lowers N. A good answer may add that outliers falling inside the band can make $\hat{e}$ a little too small, but the effect is small.