Identifying online display of Food Hygiene Rating Scheme ratings: Executive summary

Background and aims

The Food Standards Agency (FSA) wish to make display of Food Hygiene Rating Scheme (FHRS) rating mandatory for businesses online, extending the current physical display requirements in Wales and Northern Ireland and making both compulsory in England. This project was undertaken to support the rollout and enforcement of mandatory online display, by providing insight into current practice.

Approach

We created an automated solution to finding business websites and establishing whether they are displaying an FHRS rating. Data from Google Places was matched against a sample of businesses from the FHRS open data. This allowed us to collect the businesses' websites where they exist. Only an establishment's own website was included; if the Google data returned a Facebook or other social media page, a chain website, or a presence on an aggregator or booking site (e.g. OpenTable, Trivago), these were excluded.

The websites were passed through a matching pipeline that fetches the website images and matches them against reference images of FHRS ratings. Each image received a score from 0 to 100, representing the confidence of the match. The highest scoring images from each website were compiled, and all images scoring higher than a particular threshold (in this case 30) were examined manually to establish whether they were ratings or not.

Findings

Just over half of the sample (54%) had a business website. We are confident that, where a business has a website, this is being correctly obtained in around 90% of cases. After passing all websites through the matching pipeline, we estimate that the prevalence of online display is around 3% of websites. Takeaways are more likely than other business types to display a rating, while pubs are less likely. We found twice as many FHRS images in England than in Wales or Northern Ireland, despite the initial sample sizes being very similar.

Standardised, high quality images are easy to detect, while variations or low image quality/blurring make it more difficult to establish a match. Our process worked in around 80% of cases; where it did not, this was usually due to unusual features in the source code of the

websites that our method had not been able to take into account.

Policy implications and scope for future work

It is hoped that the analysis carried out here can support the drafting of legislation, by providing insight into the prevalence and nature of online display as it is now. The process can also be re-run quickly and easily, for example just prior to introducing a new policy, or in the future to gauge compliance across or for particular local authorities. The technology developed here could also form the basis of a more interactive tool to aid enforcement.

There is limited scope to improve the success rate of the process itself; there will always be variation in websites that cannot be accounted for in full, and some websites (up to 5%) are simply not accessible with programmatic methods such as these. However there is a possibility of increasing the success with which the scraper is able to grab images, and reduce the 20% failure rate by a few percent. A more standardised approach to displaying images on the part of the food businesses themselves would also improve the accuracy with which the image matching algorithm is able to identify ratings images and eliminate similar but non-matching ones.