

Identifying online display of Food Hygiene Rating Scheme ratings: Conclusion

Options for future iterations

There is some scope to reduce the proportion of websites that do not successfully pass through the matching process, which is currently around 20%, but it is limited. The scraper is looking for images in the website's source code, and expects to find these expressed in a particular way; however in around 7% of cases it could not find the images because they were represented in a slightly different way. Although it is not possible to incorporate every variation in this respect, some adjustment to the scraper could probably halve this. This would increase the success rate of the matching pipeline by a few percent.

However, beyond this, there is little from a technical point of view that could improve the scraping and matching process. There will always be quirks in websites that will fall foul of the assumptions that need to be made to enable this type of bulk, automated processing. Some websites use types of images that are not compatible with even the latest technology available to carry out image matching. And some websites will always be designed, intentionally or unintentionally, in a way that makes it impossible to retrieve their content using code.

The technology developed in this project to find and analyse food business websites could form the basis of a more interactive tool to reduce the amount of routine work involved in monitoring compliance. However, the legal implications of doing so would need to be thoroughly evaluated; they were not here, as any such tool was beyond the scope of this project. However, any future work of this nature would need to consider whether it would be in breach of legislation such as the Regulation of Investigatory Powers Act, which limits the use of online surveillance in regulatory enforcement.

Policy implications

The biggest gains in improving this image detection process would be found if the quality and standardisation of FHRS images were improved. None of the false positive images scored higher than 75, and most scored 50 or less. By comparison, the best quality FHRS images received match scores close to 100. Yet only two of the images in the full sample were found to be using this type of top quality FHRS image. More widespread use of a standardised image, of the correct type and of a minimum size, would make it much easier to detect which sites are displaying a rating.