

Crowdsourcing scoring of immunohistochemistry images: Evaluating Performance of the Crowd and an Automated Computational Method

Humayun Irshad^{1,*}, Eun-Yeong Oh¹, Daniel Schmolze¹, Liza M. Quintana¹, Laura Collins¹, Rulla M. Tamimi^{2,3}, and Andrew H. Beck¹

¹Beth Israel Deaconess Medical Center, Harvard Medical School, Department of Pathology, Boston, 02115, USA

²Harvard School of Public Health, Department of Epidemiology, Boston, 02115, USA

³Channing Division of Network Medicine, Brigham and Women's Hospital, Boston, 02115, USA

*hirshad@bidmc.harvard.edu

ABSTRACT

The assessment of protein expression in immunohistochemistry (IHC) images provides important diagnostic, prognostic and predictive information for guiding cancer diagnosis and therapy. Manual scoring of IHC images represents a logistical challenge, as the process is labor intensive and time consuming. Since the last decade, computational methods have been developed to enable the application of quantitative methods for the analysis and interpretation of protein expression in IHC images. These methods have not yet replaced manual scoring for the assessment of IHC in the majority of diagnostic laboratories and in many large-scale research studies. An alternative approach is crowdsourcing the quantification of IHC images to an undefined crowd. The aim of this study is to quantify IHC images for labeling of ER status with two different crowdsourcing approaches, image labeling and nuclei labeling, and compare their performance with automated methods. Crowdsourcing-derived scores obtained greater concordance with the pathologist interpretations for both image labeling and nuclei labeling tasks (83% and 87%), as compared to the pathologist concordance achieved by the automated method (81%) on 5,483 TMA images from 1,909 breast cancer patients. This analysis shows that crowdsourcing the scoring of protein expression in IHC images is a promising new approach for large scale cancer molecular pathology studies.

Introduction

Immunohistochemistry (IHC) is widely used for measuring the presence and location of protein expression in tissues. The assessment of protein expression by IHC provides important diagnostic, prognostic and predictive information for guiding cancer diagnosis and therapy. In the research setting, IHC is frequently evaluated using tissue microarray (TMA) technology, in which small cores of tissue from hundreds of patients are arrayed on a glass slide, enabling the efficient evaluation of biomarker expression across large numbers of patients.

The manual pathological scoring of large numbers of TMAs represents a logistical challenge, as the process is labor intensive and time consuming. Over the past decade, computational methods have been developed to enable the application of quantitative methods for the analysis and interpretation of IHC-stained histopathological images^{1,2}. While some automated methods have shown high levels of accuracy for IHC markers³⁻⁶, automated analysis has not yet replaced manual scoring for the assessment of IHC in the majority of diagnostic pathology laboratories and in many large-scale research studies.

In this study, we evaluate the use of crowdsourcing to outsource the task of scoring IHC labeled TMAs to a large crowd of users not previously trained in pathology. Over the last decade, crowdsourcing has been used in a wide range of domains, including astronomy⁷, zoology⁸⁻¹⁰, medical microbiology¹¹, and neuroscience¹²⁻¹⁴, to achieve tasks that required large-scale human labeling, which would be difficult or impossible to achieve effectively using only computational methods or domain experts.

In a pilot study, we explored the use of crowdsourcing for rapidly obtaining annotations for two core tasks in computational pathology: nucleus detection and segmentation¹⁵. This study concluded that aggregating multiple annotations from a crowd to obtain a consensus annotation could be used effectively to generate large-scale human annotated datasets for nuclei detection and segmentation in histopathological images. Crowdsourcing has also recently been evaluated for immunohistochemistry studies. Mea et al. crowdsourced 13 IHC images for detection of positive and negative nuclei and reported 0.95 Spearman correlation between pathologist and crowdsourced positivity percentages¹⁶. Recently, the Cell Slider project [CellSlider](#) by Cancer Research UK provided an online interface for members of the general public to score IHC stained TMA images, and they

reported high levels of concordance of crowdsourced scores obtained from non-experts and the scores of trained pathologists¹⁷.

The purpose of the present study is two-fold. First, we aim to evaluate the performance of crowdsourcing vs. an automated method for scoring protein expression in IHC stained TMA images. Second, we aim to evaluate the time, cost, and accuracy of two different approaches to crowdsourcing the IHC task (image-level labels vs. nucleus-level labels).

Methods

Dataset

This study used IHC-stained TMA images of breast cancer tissue from the Nurses' Health Study (NHS). The dataset consists of 5,483 scanned images of TMA cores, which were immunostained for estrogen receptor (ER) and scanned using Aperio Slide Scanner at 20× magnification. The average image size is 828 × 848 pixels. These images are derived from 1909 patients, each of whom contributed 1 – 3 TMA images, with more than half of the patients contributing 3 TMA images. All study images were scored by an expert breast pathologist, using three labels (negative=0, low positive=1 and positive=2)¹⁸.

Crowdsourcing Platform

We employed the CrowdFlower platform to design both crowdsourcing applications (image labeling and nuclei labeling). CrowdFlower is a crowdsourcing platform that works with over 50 labor channel partners to enable access to a network of more than 5 million contributors worldwide. This platform offers a number of features to improve the likelihood of obtaining high-quality work from contributors. In CrowdFlower, the job designer creates a job in the form of tasks, which are served to contributors for labeling. Each task is a collection of one or more images sampled from the data set. The job designer creates test questions (test images which have been previously labeled by pathologists) that are used for dual purposes: qualification of contributors during quiz mode and monitoring of contributors during judgment mode. Contributors must maintain a defined level of accuracy on the test questions to be permitted to complete the job. In addition, the job designer specifies the payment per task and the number of labels desired per image. After job completion, CrowdFlower provides a list of labels (annotations) for all the images. Additional information on the CrowdFlower platform is available at www.crowdflower.com.

Job Design and Crowdsourcing Applications

Each crowdsourcing job has two modes: quiz mode and judgment mode. Quiz mode occurs at the beginning of a job. In quiz mode, there is only one task and the task consists of 5 test question images. In judgment mode, there are a number of tasks and each task consists of 4 actual images and one test image which is presented to the contributor in the same manner as the unlabeled images such that the contributor is unaware if he/she is annotating an unlabeled image or test image. Each contributor must qualify during quiz mode to enter in judgment mode and can remain in judgment mode as long as his/her accuracy on test questions is above a threshold level. For ensuring high quality of labels, we defined five parameters which may influence labeling performance.

- The first is test question minimum accuracy that ensures each contributor must maintain minimum 60% accuracy on test questions throughout the job completion.
- The second is minimum time per task that ensures each contributor must spend a minimum of 10 seconds to complete one task.
- The third is maximum number of judgments per contributor that enable more contributors to participate in the job. In our jobs, we defined maximum number of judgment per contributor 500 judgments.
- The fourth is a minimum number of images (20) for the contributor to review in work mode prior to computing a trust score for each worker and prior to filtering workers based on their trust score.
- The fifth is the number of labels to collect per image. We collect three labels per image for both jobs.

Our study includes two types of labeling jobs: image labeling and nuclei labeling. Figure 1 illustrates the flow chart of both crowdsourcing jobs. Each job contains instructions, which provide examples of expert-derived labels and guidance to assist the contributor in learning the process of labeling.

Image Labeling

In the image labeling job, each contributor estimates the percentage of cancer nuclei stained brown (positive) and blue (negative) in the image and then selects the image label (score) depending on given criteria: if percentage of brown nuclei is less than 1% then image label is A (negative protein expression), if percentage is between 1% and 10%, then image label is B (low positive protein expression), if percentage is between 10% and 50% then image label is C (positive protein expression) and if percentage is more than 50% then image label is D (high positive protein expression). The total pool of test question images used in both quiz and judgment modes are 250, which are labeled by pathologists. Figure 2 shows the interface for image labeling.

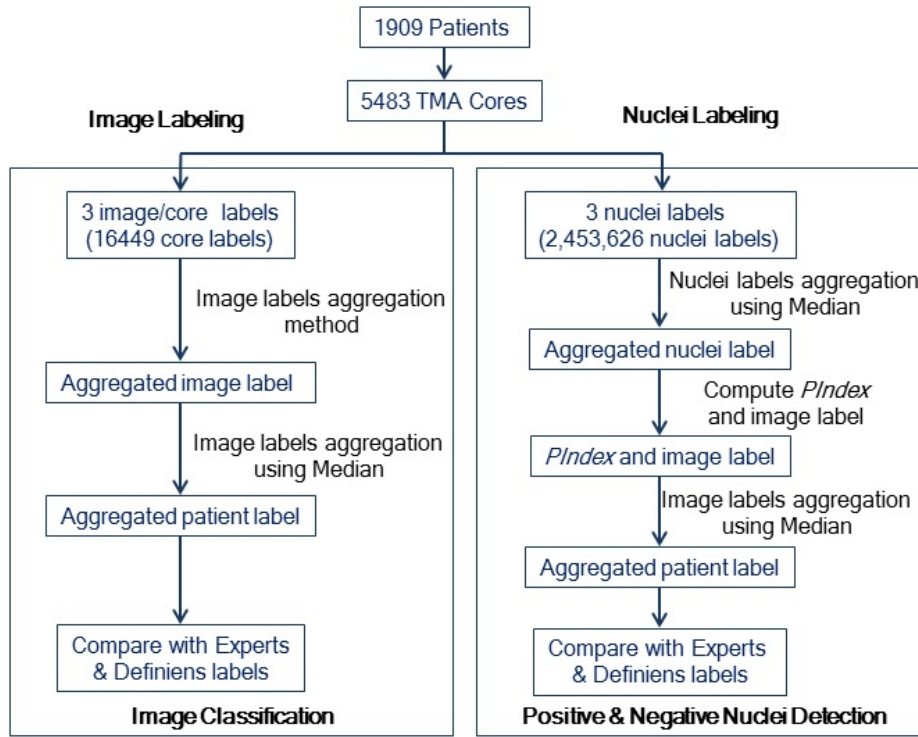


Figure 1. Crowdsourcing work flow for Image Labeling and Nuclei Labeling.

Nuclei Labeling

In the nuclei labeling job, we ask contributors to detect positive and negative nuclei in the image. In the nuclei labeling job, we first ask contributors to identify the presence of nuclei in the image (yes/no). If they do identify the presence of nuclei, then we ask the contributor to label the nuclei using a dot operator (by clicking at the center of each nucleus). At completion of job, CrowdFlower provides the position of positive and negative nuclei in the images. For each image, we collected positive and negative nuclei from three different contributors. The total pool of test question images used in both quiz and judgment modes for nuclei labeling are 100 images, which are labeled by pathologists. After counting number of positive and negative nuclei, we compute the positivity index ($PIndex = \frac{No.ofPositiveNuclei}{TotalNo.ofNuclei}$). From positivity index, we compute the image label using image labeling criteria (mentioned in Image Labeling section). Figure 3 shows the interface for nuclei labeling.

Aggregation Methods for Image Labeling Problem

We calculated the aggregated label for each image using four different methods: maximum crowd votes (CV), maximum crowd trust scores (CT), maximum weighted crowd votes (ωCV) and maximum weighted crowd trust scores (ωCT). CV is computed by summing the votes for each label and selecting the label with the maximum number of votes as the aggregated label. CT is computed by summing the contributor trust score (CT) for each label and selecting the label with the maximum trust score as the aggregated label. For ωCV and ωCT methods, we multiply the class weights with crowd votes for each label and crowd trust scores for each label, respectively.

$$WeightedCrowdVote = \frac{\omega_A V_A + \omega_B V_B + \omega_C V_C + \omega_D V_D}{V_A + V_B + V_C + V_D}$$

$$WeightedCrowdTrustScore = \frac{\omega_A T_A + \omega_B T_B + \omega_C T_C + \omega_D T_D}{T_A + T_B + T_C + T_D}$$

where V_A, V_B, V_C and V_D are crowd votes from each class labels; T_A, T_B, T_C and T_D are sum of crowd trust scores for each class labels; and $\omega_A, \omega_B, \omega_C$ and ω_D are class weights. We calculated the class weights by taking the mean of lower and upper boundary of the class. For class A, lower boundary is 0 and upper boundary is 0.01, the weight of class A is 0.005. For class B, lower boundary is 0.01 and upper boundary is 0.1, the weight of class B is 0.05. For class C, lower boundary is 0.1 and upper boundary is 0.5, the weight of class C is 0.3. For class D, lower boundary is 0.5 and upper boundary is 1, the weight of class D

Estrogen Receptor Protein Expression In Breast Cancer - Image Labeling

Instructions ▾




Image label based on percentage of brown nuclei:

- Class A - Percentage of nuclei that are brown is less than 1%.
- Class B - Percentage of nuclei that are brown is between 1% and 10%.
- Class C - Percentage of nuclei that are brown is between 10% and 50%.
- Class D - Percentage of nuclei that are brown is more than 50%.
- Class E - Unable to interpret (blurry image, too few nuclei (less than 20), or other).

Figure 2. Crowdsourcing Application interface for image labeling. The screenshot illustrates the interface for selecting the image class label.

is 0.75. The aggregated label selected is the label whose class bounds contain the weighted crowd vote or weighted crowd trust score.

Sensitivity Analysis for Different Combinations of Crowd Size

To estimate the number of crowd labels required to generate optimal aggregated crowd label, we performed a sensitivity analysis of aggregated labels using different combination of crowd sizes. For this pilot study, we collected 10 crowd labels for each image, and we computed the aggregated label of each image using different combination of crowd sizes (1 to 10), according to Algorithm 1.

Performance Measures

We explored different performance measures to evaluate inter-observer reliability of scores. For measuring the inter-observer reliability, we measured percent agreement (A_g) or accuracy, which is calculated as the number of agreed labels divided by total number of labels, Kappa (κ) which measures the agreement among observers adjusted for the possibility of by chance agreement, Spearman correlation (ρ) which measures the mean of bivariate Spearman's rank correlations between observers for inter-observer reliability, and intra-class correlation (ICC). For image classification, we used A_g performance measures for comparing different methods of label aggregation and the automated method.

Estrogen Receptor Protein Expression In Breast Cancer - Labeling Of Negative And Positive Cells Using Crowdsourcing



Figure 3. Crowdsourcing application interface for nuclei labeling. The screenshot illustrates the interface for labeling the positive and negative nuclei separately.

Results

Image labeling on 380 TMA cores - A pilot study

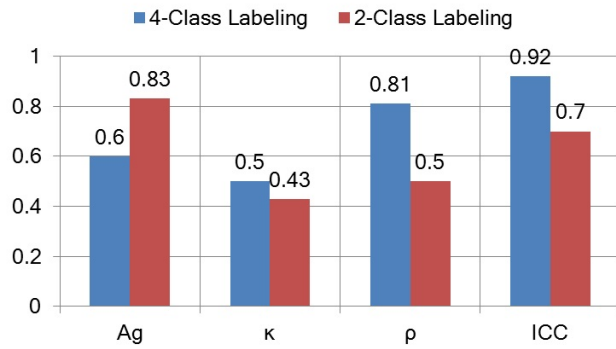
We designed a pilot study to test the crowd sourcing application for IHC image labeling and to assess the improvement in crowdsourcing performance as we increase the numbers of aggregated instances per image. In the pilot study, we collected 10 crowdsourced labels for 380 images. We also collected three pathologist labels for each of these 380 images using the same crowd sourcing interface.

We assessed inter-observer reliability among pathologists using 4-class labeling as well as 2-class labeling as shown in Figure 4(a). For 2-class labeling, we merged all positive classes (B, C and D) into a single positive class (B). We observed Kappa values of 0.43 and 0.5 for 2-class and 4-class labeling, respectively, indicating moderate inter-pathologist agreement in IHC interpretation.

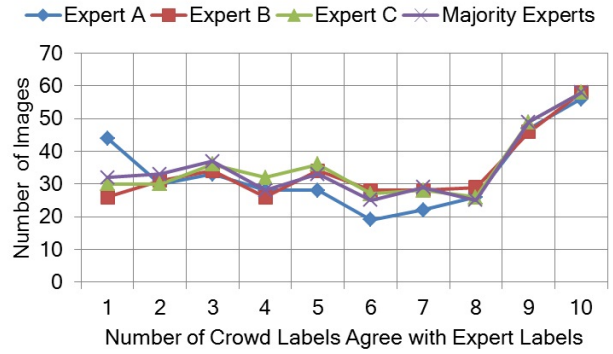
For 380 images, we obtained 10 crowd labels per image and compared these scores with the pathologist scores as shown in figure 4(b). We found a wide range of agreement on images between the crowd and the pathologist scores, with a median level

Algorithm 1 Sensitivity Analysis for Different Crowd Sizes

```
for all Crowd Sizes:  $C_i$  in  $i \in 1, 2, 3, \dots, 10$  do  
   $P =$  Compute combination patterns (without replacement) for Crowd Size  $C_i$   
  for all Pattern:  $P_j$  in  $j \in 1, 2, 3, \dots, J$  do  
    for all Images:  $I_k$  in  $k \in 1, 2, 3, \dots, K$  do  
      Compute the Aggregated labels for combination pattern  $P_j$  of Crowd Size  $C_i$  for Image  $I_k$   
    end for  
  end for  
  Compute Agreement of Aggregated Labels with GT Labels  
end for
```

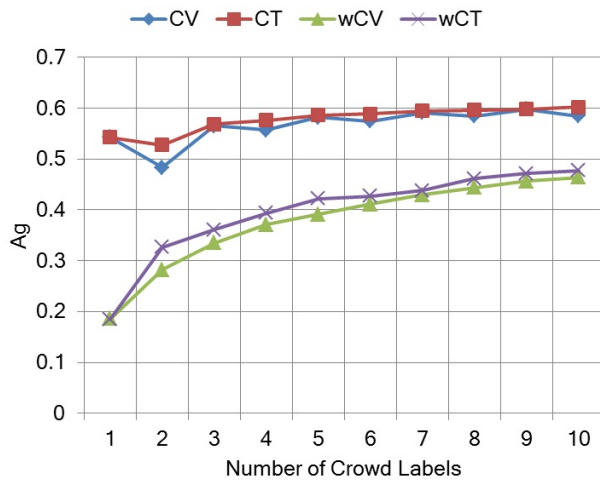


(a) Inter-observer reliability of pathologist labels

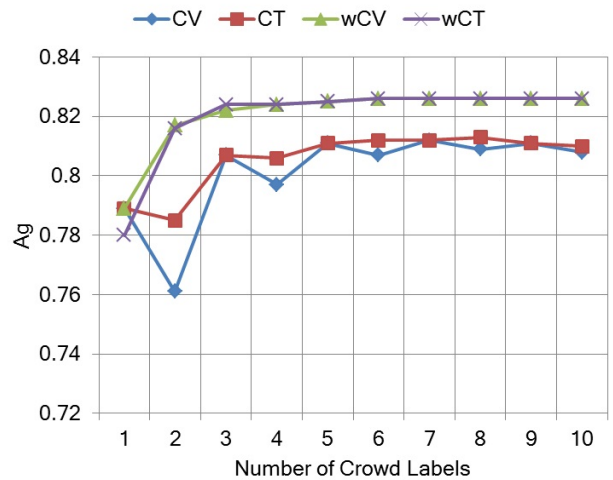


(b) Number of crowd labels agree with pathologist labels

Figure 4. Inter-observer reliability of pathologist labels and agreement with crowd labels.



(a) 4-class labeling sensitivity analysis



(b) 2-class labeling sensitivity analysis

Figure 5. Sensitivity analysis of crowd labels in pilot study. The analysis supports using 3 crowdsourced labels per image.

of agreement of 6/10 and 15% of images showing 10/10 agreement. For a range of crowd labels per image (ranging from 1 to 10), we computed aggregated labels and assessed the agreement of the consensus score with the pathologist score for each number of crowd labels as reported in Figures 6(a) and 6(c). The A_g is not significantly improved after crowd size 3 for 4-class and 2-class image labeling problem.

Image labeling on 5483 TMA cores

Based on the results of the sensitivity analysis, we collected 3 crowd labels for each image for the image and nucleus labeling crowdsourcing jobs for the main study. The image and nuclei labeling work flow is illustrated in Figure 1. For 4-class image labeling, we collected 3 labels for each TMA image using CrowdFlower as shown in Figure 2. In total, 16,449 image labels were collected for 5,483 images. Aggregated image labels were computed with four aggregation methods (CV, CT, ω CV and ω CT). The aggregated label at patient level was computed by taking the median of all aggregated image labels belonging to that patient. Pathologists labeled these images using 3-class labeling: negative (A), low positive (B) and positive (C). Since we obtained 4-class labeling from the crowd, to compare the crowd labels with pathologist labels, we merged crowd class D into crowd class C. The CV aggregation method reported higher A_g and Spearman ρ than other aggregation methods for 3-class labeling as reported in Table 1. For 2-class labeling, we merged all positive classes (B and C) into a single positive class (B) for both crowd and pathologist aggregated labels. The CV aggregation method outperformed as compared to other aggregation method for A_g and ρ as reported in Table 1.

Crowdsourcing Types	Method	3-Class Labeling		2-Class Labeling	
		A_g	ρ	A_g	ρ
Image Labeling	Crowd CV	0.71	0.64	0.83	0.62
	Crowd CT	0.68	0.65	0.81	0.61
	Crowd ω CV	0.64	0.63	0.77	0.59
	Crowd ω CT	0.64	0.64	0.77	0.59
Nuclei Labeling	Crowd	0.77	0.68	0.87	0.63
	Definiens	0.70	0.51	0.81	0.48

Table 1. Crowdsourced image labeling and nuclei labeling results.

Nuclei Labeling on 5483 TMA cores

For the nuclei labeling job, we collected 3 nuclei labels for all 5,483 TMA images. Total number of nuclei labels was 2,453,646. The aggregated number of positive and negative nuclei was calculated for each image as the median number of positive and negative nuclei labeled by the crowd. Then, we computed $PIndex$ for each image. $PIndex$ was converted into image labels (A,B and C) according to following image labeling criteria:

$$\text{Image Label} = \begin{cases} \text{Class C (positive)} & \text{if } PIndex > 0.1 \\ \text{Class B (low positive)} & \text{if } PIndex > 0.01 \text{ and } PIndex \leq 0.1 \\ \text{Class A (negative)} & \text{if } PIndex \leq 0.01 \end{cases}$$

Lastly, we computed the aggregated patient label by taking the median of the all the image labels belonging to that patient and compared with the pathologist labels as reported in Table 1. We also performed 2-class labeling by merging all positive classes into a single positive class for both crowd and pathologist labels. The A_g and ρ are 0.77 and 0.68 for 3-class labeling and 0.87 and 0.63 for 2-class labeling, respectively.

In order to compare with an automated method, we developed an image processing pipeline in Definiens Tissue Studio. This pipeline detected positive and negative nuclei in TMA images and computed the $PIndex$. The crowdsourcing $PIndex$ was correlated with the Definiens $PIndex$ (ρ is 0.75). However, considering pathologist labels as ground truth, both types of crowdsourcing jobs (image and nuclei labeling) resulted in higher A_g and ρ than Definiens for both 3-class labeling and 2-class labeling.

The Crowd showed significantly improved performance on test questions for the nuclei labeling job as compared with the image labeling task as reported in Table 2. This finding supports the overall higher level of accuracy seen with the nuclei labeling approach as compared with the image labeling approach.

Crowdsourcing Performance

We first assessed the contributor (crowd) performance for both crowdsourcing jobs. The number of contributors who participated in both jobs is shown in Table 2. The contributors who maintained the minimum accuracy (60%) on test questions during quiz and work modes are trusted contributors and the rest are untrusted contributors. In work mode, there were 61 trusted contributors for image labeling and 2,216 for nuclei labeling. The average time of trusted contributors was 32 seconds for image labeling and 306 seconds for nuclei labeling per image while the average time of untrusted contributors was 149 seconds for image labeling and 207 seconds for nuclei labeling. Thus, trusted contributors took less time to label images as compared to untrusted contributors; however, trusted contributors took more time to label nuclei as compared to untrusted contributors. These results suggest that efficient labeling of nuclei is a complex job requiring sufficient time for strong performance. Figure 6 illustrates the distribution of crowd trust scores for both jobs. The image labeling contributors have higher trust score as compared to nuclei labeling contributors. The average test question accuracy for trusted contributors is 80% for image labeling and 76% for nuclei labeling while average test question accuracy for untrusted contributors is 66% for image labeling and 42%

Crowdsourcing Jobs	Quiz Mode		Work Mode	
	Passed	Failed	Passed	Failed
Image Labeling	113	155	61	52
Nuclei Labeling	3244	1572	2216	1243

Table 2. Crowd performance on test questions in quiz mode and work mode.

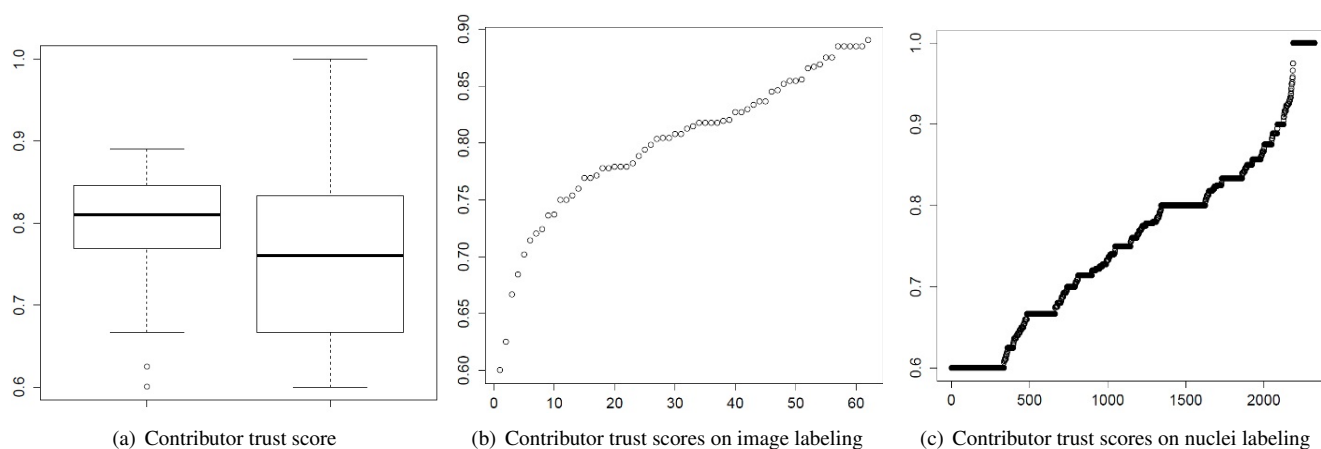


Figure 6. Contributor trust scores analysis on Crowdsourcing jobs.

for nuclei labeling. The trust scores were moderately correlated with the number of images labeled; $\rho = 0.41$, $P < 0.0008$ for image labeling job and $\rho = 0.186$, $P < 2.2e^{-16}$ for nuclei labeling. The average time for image labeling is 50 seconds and nuclei labeling per image is 373 seconds.

The image labeling job was finished in 4 hours and nuclei labeling was finished in 472 hours. The Crowdfunder platform charged \$ 282 for image labeling and \$ 2,280 for nuclei labeling job. These data suggest that although nuclei labeling produced some improvements in accuracy, it cost significantly longer in terms of time to complete the full job (118 fold longer) and cost (8 fold more expensive).

Discussion

The principle of applying crowdsourcing in science, which has enabled whale sound classification¹⁰, malaria parasite classification¹¹, sleep spindle detection¹³ and nuclei detection and segmentation in histopathology¹⁵, has become increasingly well established in recent years. Crowdsourced work can be used to classify objects (whale sound and malaria parasite classification), detect objects (nuclei detection) and segment objects (nuclei segmentation). The aim of this study was to better understand how to use crowdsourcing for IHC image interpretation. This laborious and time consuming image quantification task has also been performed using automated methods^{4-6, 19-21}. However, no prior studies have directly compared crowdsourcing vs. automated methods in the interpretation of IHC.

In this study, we quantify IHC TMA images for labeling of ER status with two different crowdsourcing approaches, image labeling and nuclei labeling. In the image labeling task, the crowd was asked to estimate the percentage of positive cells for each IHC image, while in nuclei labeling task, the crowd was asked to label individual nuclei within IHC images as either positive or negative. We completed these crowdsourcing tasks on a large data set containing 5483 TMA images belonging to 1909 patients, which were previously labeled by an expert pathologist and by an automated method. In our study, crowdsourcing-derived scores obtained greater concordance with the pathologist interpretation for both image labeling and nuclei labeling tasks, as compared with the pathologist concordance achieved by the automated method.

Overall, the crowdsourced scores produced from nuclei labeling (as opposed to image labeling) showed somewhat higher agreement with the pathologist scores; however, the time and cost required for the nuclei labeling far exceeded the time and cost for the image labeling. Nuclei labeling is more laborious task spread over many people, even though paying more and takes longer time for nuclei scoring, still cheaper and quicker than using pathologists. Our study results support that crowdsourcing is a promising new approach for scoring biomarker studies in large scale cancer molecular pathology studies. A limitation of our current crowdsourcing application is that we do not ask the Crowd to classify nuclei into specific types (e.g., cancer epithelial nucleus, lymphocyte nucleus). We expect the addition of training the crowd to classify cell types in addition to classifying IHC positivity will further improve crowd performance, although the incorporation of cell type-specific scoring may increase the time and cost of the overall task. This represents an important direction for future research.

References

1. Gurcan, M. N. *et al.* Histopathological image analysis: a review. *Biomedical Engineering, IEEE Reviews in* **2**, 147–171 (2009).

2. Irshad, H., Veillard, A., Roux, L. & Racoceanu, D. Methods for nuclei detection, segmentation, and classification in digital histopathology: A review—current status and future potential. *Biomedical Engineering, IEEE Reviews in* **7**, 97–114 (2014).
3. Giltneane, J. M. & Rimm, D. L. Technology insight: identification of biomarkers with tissue microarray technology. *Nature clinical practice Oncology* **1**, 104–111 (2004).
4. Bolton, K. L. *et al.* Assessment of automated image analysis of breast cancer tissue microarrays for epidemiologic studies. *Cancer Epidemiology Biomarkers & Prevention* **19**, 992–999 (2010).
5. Ali, H. *et al.* Astronomical algorithms for automated analysis of tissue protein expression in breast cancer. *British journal of cancer* **108**, 602–612 (2013).
6. Howat, W. J. *et al.* Performance of automated scoring of er, pr, her2, ck5 and egfr in breast cancer tissue microarrays in the breast cancer association consortium. *The Journal of Pathology: Clinical Research* **1**, 18–32 (2015).
7. Lintott, C. J. *et al.* Galaxy zoo: morphologies derived from visual inspection of galaxies from the sloan digital sky survey. *Monthly Notices of the Royal Astronomical Society* **389**, 1179–1189 (2008).
8. Sullivan, B. L. *et al.* ebird: A citizen-based bird observation network in the biological sciences. *Biological Conservation* **142**, 2282–2292 (2009).
9. Marris, E. Supercomputing for the birds. *Nature* **466**, 807–807 (2010).
10. Shamir, L. *et al.* Classification of large acoustic datasets using machine learning and crowdsourcing: Application to whale calls. *The Journal of the Acoustical Society of America* **135**, 953–962 (2014).
11. Luengo-Oroz, M. A., Arranz, A. & Frean, J. Crowdsourcing malaria parasite quantification: an online game for analyzing images of infected thick blood smears. *Journal of medical Internet research* **14**, e167 (2012).
12. Kim, J. S. *et al.* Space-time wiring specificity supports direction selectivity in the retina. *Nature* **509**, 331 (2014).
13. Warby, S. C. *et al.* Sleep-spindle detection: crowdsourcing and evaluating performance of experts, non-experts and automated methods. *Nature methods* **11**, 385–392 (2014).
14. Arganda-Carreras, I. *et al.* Crowdsourcing the creation of image segmentation algorithms for connectomics. *Frontiers in neuroanatomy* **9** (2015).
15. Irshad, H. *et al.* Crowdsourcing image annotation for nucleus detection and segmentation in computational pathology: Evaluating experts, automated methods, and the crowd. In *Pacific Symposium on Biocomputing (PSB)* (2015).
16. Della Mea, V., Maddalena, E., Mizzaro, S., Machin, P. & Beltrami, C. A. Preliminary results from a crowdsourcing experiment in immunohistochemistry. *Diagnostic pathology* **9**, S6 (2014).
17. dos Reis, F. J. C. *et al.* Crowdsourcing the general public for large scale molecular pathology studies in cancer. *EBioMedicine* **2**, 679–687 (2015).
18. Collins, L. C., Marotti, J. D., Baer, H. J. & Tamimi, R. M. Comparison of estrogen receptor results from pathology reports with results from central laboratory testing. *Journal of the National Cancer Institute* **100**, 218–221 (2008).
19. Mohammed, Z. *et al.* Comparison of visual and automated assessment of ki-67 proliferative activity and their impact on outcome in primary operable invasive ductal breast cancer. *British journal of cancer* **106**, 383–388 (2012).
20. Inwald, E. *et al.* Ki-67 is a prognostic parameter in breast cancer patients: results of a large population-based cohort of a cancer registry. *Breast cancer research and treatment* **139**, 539–552 (2013).
21. Gudlaugsson, E. *et al.* Comparison of the effect of different techniques for measurement of ki67 proliferation on reproducibility and prognosis prediction accuracy in breast cancer. *Histopathology* **61**, 1134–1144 (2012).

Acknowledgements

The research reported in this publication was supported in part by the National Library of Medicine of the National Institutes of Health under Award Number K22LM011931. We would like to thank the participants and staff of the Nurses' Health Study (NHS) for their valuable contributions as well as the following state cancer registries for their help: AL, AZ, AR, CA, CO, CT, DE, FL, GA, ID, IL, IN, IA, KY, LA, ME, MD, MA, MI, NE, NH, NJ, NY, NC, ND, OH, OK, OR, PA, RI, SC, TN, TX, VA, WA, WY. The authors assume full responsibility for analyses and interpretation of these data. The data collection and creation of TMAs in this publication was supported by Public Health Service Grant under Award Number CA087969 and National Cancer Institute of National Institutes of Health under Award Number CA186107. This investigation was approved by the institutional Review Board at the Brigham and Women's Hospital and the Harvard School of Public Health.

Author contributions statement

H.I. and A.H.B. designed the study. H.I. developed crowdsourcing application, performed experiments and analyzed the results. E.O., D.S., L.M.Q and L.C. are pathologists and performed image labeling. R.M.T. provides NHS data. All authors reviewed the manuscript.

Competing Financial Interests

The author(s) declare no competing financial interests.