

Gender and Performance in Computer Science

ISABEL WAGNER, University of Hull

The term *gender gap* refers to the significant underrepresentation of females in many subjects. In Computer Science, the gender gap exists at all career levels. In this paper, we study whether there is a performance gap in addition to the gender gap. To answer this question, we analyzed statistical data on student performance in Computer Science from 129 universities in the UK covering the years 2002 to 2013. We find that male students were awarded significantly more first-class degrees than female students. We evaluate four other subjects – Subjects Allied to Medicine, Business & Administrative Studies, Mathematical Sciences, and Engineering & Technology – and find that they do not exhibit this performance gap. Following on from this finding, we review explanations for the gender and performance gaps, as well as potential solutions to eliminate the gaps. Most solutions do not require major institutional change and could thus be implemented easily.

CCS Concepts: • **Social and professional topics** → **Computing education programs; Women;**

Additional Key Words and Phrases: Gender gap, performance gap

ACM Reference Format:

Isabel Wagner. 2015. Gender and Performance in Computer Science. *ACM Trans. Comput. Educ.* V, N, Article A (January YYYY), 17 pages.

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

Less than 16% of Computer Science (CS) degrees in the UK were awarded to women in 2013, according to data from the Higher Education Statistics Agency (HESA). This significant underrepresentation of females in CS is commonly referred to as the *gender gap*. Interestingly, the gender gap exists in some countries, but not others: for example, it exists in the United States [Fisher et al. 1997], Greece [Papastergiou 2008], and Italy [Boschetto and Cortesi 2009], but does not exist in India, Armenia [Gharibyan and Gunsaulus 2006], Malaysia [Mellström 2009], Serbia [Ivanović et al. 2010], and Vietnam [Shillabeer and Jackson 2013]. Many authors have emphasized the importance of eliminating the gender gap, citing, among other reasons, benefits for the IT industry [Ashcraft and Blithe 2009], innovation, and productivity [Prey and Weaver 2013].

In this study, we focus on another aspect of the gender gap: the performance gap, or systematic difference in the performance of male and female students. If a performance gap exists, it would be a factor that discourages females to take up computer science as their profession. Here, we analyze statistical data from a large number of UK universities to determine the magnitude of the performance gap in CS. Using the results of this study, we can adjust our efforts towards recruiting and supporting female students, and thus achieve a long-term increase in the number of women in computer science. To the best of our knowledge, only small-scale studies have so far been performed to investigate whether the performance gap exists.

Author's addresses: (Current address) I. Wagner, De Montfort University, The Gateway, Leicester, LE1 9BH, United Kingdom

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© YYYY ACM. 1946-6226/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

2. RESEARCH QUESTIONS

Many studies report a difficult classroom environment for female students that is caused, at least in part, by females being in a minority [Margolis et al. 2000; Lewis et al. 2006]. However, it is unclear whether this difficult environment has a negative impact on female classroom performance. One may argue to the contrary: because CS is not a mainstream choice for women, those who do choose it are more committed and thus perform better [Barker et al. 2009; Ivanović et al. 2010]. Other studies report no significant performance differences [Wilson 2002; Byrne and Lyons 2001]. The discrepancies between these studies might be explained by the fact that they were small-scale studies relying on data from only one institution or even a single cohort.

In this paper, we conducted a large-scale analysis of student performance using statistical data spanning 12 years and 129 universities in the UK. Using this data, we investigated how the percentage, number, and gender of students in a CS program correlate with their performance. Our 3 hypotheses were that:

- Female computer science students perform worse than male students (hypothesis A);
- Female students perform better as their total number and percentage increases (hypothesis B);
- Both male and female students perform better as the year of entry increases (hypothesis C).

3. METHODOLOGY

We use multiple regression to analyze statistical data from the publicly available Higher Education Information Database for Institutions (*heidi*)¹ to gather evidence on hypotheses A, B, and C. *Heidi* data is available for a period of 12 years, from 2002 to 2013, and for a total of 204 higher education institutions in the UK. Of these, only 129 institutions were used, resulting in a total of 1372 data points (unique institution/year combinations); the others did not award CS degrees in some or all years. Each data point represents between 5 and 520 students (165 students on average).

3.1. Result variables: Performance indicators

Two variables available in *heidi* can be used to indicate student performance: the degree classification obtained, and the employment status after graduation. Employment status is a problematic indicator for two reasons. First, employment status in *heidi* is not based on university-reported numbers, but rather on self-reported data, which may introduce bias. Second, whether a student finds employment or not depends on many other factors besides their performance [Gibbs 2010]. Thus, we based our study on the degree classification variable.

In the UK, students can earn one of four degree classes, depending on their performance: a first class degree (an average of 70% or higher across their assessed work); an upper second class, or 2.1 (60% or higher); a lower second class, or 2.2 (50% or higher); and a third class degree (40% or higher). Two indicators are commonly used to evaluate the performance of a group of students: the proportion of students who obtain a first class degree, and the proportion of students who obtain a ‘good’, i.e. a first or upper second class degree. This study investigates both indicators (referred to from here as ‘first’ and ‘good’).

¹<https://heidi.hesa.ac.uk/>

3.2. Explanatory variables

In this study, we consider how student gender, the percentage of female students in a CS program², and the total number of female students in a CS program correlate with the performance indicators. We also consider graduation year [Yorke 2009] and UCAS entry tariff [McKenzie and Schweitzer 2001], which have previously been shown to be correlated with student performance.

In the UK, the UCAS entry tariff indicates how good a student's qualifications for university entry are. To determine a student's entry tariff, the Universities and Colleges Admissions Service (UCAS) assigns points to pre-university qualifications. The sum of a student's points constitute the entry tariff, which is frequently used as a university entry requirement and ranges from 0 (worst) to 768 (best) points. All subjects combined, the average female student had 15 more UCAS points than the average male (2013 data, 351 vs. 336 points). The point distance was the same for Computer Science, with a slightly higher average (393 vs. 378 points).

3.3. Data and data quality

Two *heidi* tables provided the data for this study. First, the *HE qualifiers* table reports the number of CS students achieving each degree classification, grouped by gender. Second, the *UCAS accepted applicants* table reports the number of students each institution accepts into their Computer Science and Maths programs, grouped by gender and UCAS tariff.

Two key characteristics of *heidi* data might influence data quality: rounding and consistency. Three consistency issues arise. First, institutions may report data incorrectly. Second, the UCAS tariff may not reflect the true entry qualifications of a student because it does not weight subject-specific qualifications higher than general qualifications. For example, maths proficiency is one predictor of success in CS [Konvalina et al. 1983; Newhall et al. 2014], but the UCAS tariff does not give special weight to maths qualifications. Third, UCAS tariff and degree classification cannot be analyzed on a per-student basis because they are reported in separate tables. In addition, *heidi* does not provide data on students who dropped out or switched subjects. Thus UCAS tariff and degree classification can only be analyzed on a per-institution level, and the group of students receiving degrees can be different from the group that entered.

Rounding is a practice in which student numbers are rounded to the nearest multiple of 5 to protect student privacy. Rounding was problematic for our study, because low numbers of female CS students lead to artifacts, especially when a single institution is considered. However, in an analysis considering many institutions rounding tends to average out.

In addition to the explanatory variables considered in this study, the *heidi* database contains data for several other variables shown to correlate with performance, such as age and socio-economic background [McKenzie and Schweitzer 2001]. However, we found that slicing the data according to more variables reduced the numbers within each slice, and therefore exacerbated rounding artifacts. Therefore, we disregarded the effects of age and socio-economic background.

4. RESULTS

We used multiple regression to analyze the correlation between the 'first' and 'good' performance indicators (result variables), and the explanatory variables gender, percentage of female students, number of female students, total number of students, graduation year, and UCAS entry tariff, as well as higher-order terms and interactions

²We use US terminology here: a *program* consists of a number of *courses* and leads to a degree. Terminology is different in the UK, where a *course* consists of a number of *modules* and leads to a degree.




Fig. 1. Correlation of gender and the proportion of first class degrees




Fig. 2. Influence of graduation year on the proportion of first class degrees, by gender

between the explanatory variables. The regression uses a generalized linear model (because result variables were proportion data), and the quasibinomial error family (because of overdispersion [Crawley 2012]).

4.1. First Class Degrees

Through a process of stepwise model simplification and elimination of non-significant variables from the model, we arrived at a model containing five significant variables that contribute to explaining variations in the proportion of first class degrees. These are the UCAS tariff ($p < 0.001$), gender ($p < 0.001$), year ($p < 0.001$), the number of females in a program ($p < 0.01$), and an interaction between UCAS tariff and gender ($p < 0.001$). None of the other interactions or higher-order terms were significant.

Gender. Male students were awarded a significantly higher percentage of first class degrees than females (17.5% vs. 15.7%, 95% confidence interval for the difference [1.3%, 2.3%], $p < 0.001$). This result is shown in Figure 1, where the boxes indicate the three quartiles, with vertical lines extending to the highest resp. lowest values within 1.5 times the inter-quartile range. The dots represent data points outside of this range. The notches indicate a 95% confidence interval for comparing medians [Mcgill et al. 1978]. In Figure 1, there is no overlap between the notches for female and male students, which indicates that the difference between the medians is significant at the 5% level. This supports hypothesis A, indicating that there is a gender gap in the proportion of first class degrees.




Fig. 3. Influence of UCAS tariff on the proportion of first class degrees awarded to students, by gender

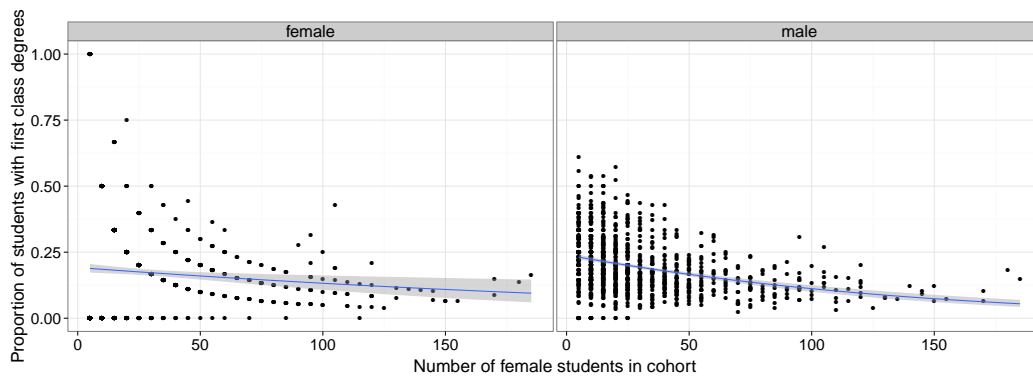


Fig. 4. Influence of the number of females on the proportion of first class degrees, by gender

Year. We then asked whether the proportion of first class degrees has increased as the year of entry increases (hypothesis C) for both male and female students. Figure 2 shows the data points as well as the regression line in a scatter plot. The shaded area around the regression line indicates the 95% confidence interval. Figure 2 indicates that the proportion of first class degrees has increased as the year of entry increases, supporting hypothesis C (a phenomenon referred to as grade inflation [Yorke 2009]).

UCAS entry tariff and interaction with gender. We then asked whether institutions requiring high entry qualifications award a higher proportion of first class degrees. Figure 3 shows that the proportion of first class degrees and the average UCAS tariff of students in each institution are positively correlated. The fact that qualification on entry is a good predictor of performance is expected, and has previously been discussed in the literature [Konvalina et al. 1983; Byrne and Lyons 2001].

Strikingly, the slope of the regression line for male students in Figure 3 is significantly steeper than the slope for females (0.0037 vs. 0.0026, the male slope is 42% steeper than the female slope, $p < 0.001$), indicating that females, especially those with high entry qualifications, are awarded first class degrees at a significantly lower rate than male students. This result provides support for hypothesis A (females perform worse).

Number of female students. We hypothesized that females would perform better when their total number in a program increases (hypothesis B). Figure 4 shows that




Fig. 5. Decrease in the proportion of female students between 2002 and 2013

an increasing number of female students in a program is correlated with decreasing performance of both males and females. This indicates that the performance of female students does not benefit from a large number of female peers, and is evidence against hypothesis B.

This finding can be explained by the “cream of the crop” effect [Milyo and Schosberg 2000], which states that an increasing percentage of females in a field leads to a decrease in average performance because the best females join the field first (for a fuller explanation, see Section 5). We can investigate this effect further by observing that the percentage of female students has decreased from 22% in 2002 to 16% in 2013 (see Figure 5). We therefore analyze two eras: the 2011-2013 era with a low percentage of females and the 2002-2004 era with a high percentage. In 2002-2004, males were awarded a significantly higher proportion of first class degrees (13.5% vs. 12%, $p < 0.001$), whereas the difference was not significant in 2011-2013 (25.7% vs. 24.7%, $p < 0.18$). This increase in female performance coupled with a decrease in the percentage of females is exactly what the cream of the crop effect predicts.

Other variables. Of the explanatory variables initially considered in the regression, two were not statistically significant. The total number of students ($p > 0.09$) has a negligible influence on the proportion first class degrees, as shown in Figure 6. The percentage of female students seems to be correlated with decreasing performance of both males and females (Figure 7), but this result is not statistically significant ($p > 0.46$). The latter is particularly interesting, because hypothesis B states that a low percentage of female students would be correlated with low performance. This is apparently not the case, and constitutes further evidence against hypothesis B.

In summary, when the proportion of first class degrees is used as performance indicator, we have found evidence to support hypotheses A (females perform worse) and C (grade inflation), and evidence against hypothesis B (females perform better when their number/percentage increases).

4.2. Good Degrees

We hypothesized that ‘good’ degrees might be affected differently from first class degrees. To investigate this, we again used multiple regression and a process of stepwise model simplification. The regression model for the ‘good’ degrees result variable contains the same five significant variables as the model for first class degrees: UCAS tariff ($p < 0.001$), gender ($p < 0.001$), year ($p < 0.001$), number of females in a program ($p < 0.001$), and an interaction between UCAS tariff and gender ($p < 0.01$).

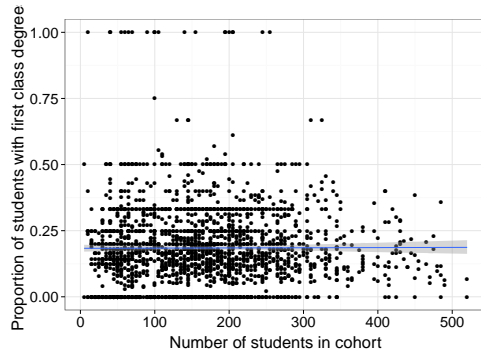


Fig. 6. Correlation between the total number of students in a cohort and the proportion of first class degrees




Fig. 7. Correlation between the percentage of female students and the proportion of first class degrees




Fig. 8. Correlation of gender with the proportion of good degrees

Gender. Figure 8 shows that the correlation between gender and the proportion of good degrees is reversed in comparison with the proportion of first class degrees: female students are awarded a higher proportion of good degrees than males (56% vs. 53.9%, 95% confidence interval for the difference $[-2.7\%, -1.5\%]$, $p < 0.001$). This is evidence against hypothesis A, indicating that there is no performance gap in the proportion of good degrees.

Year. We then asked whether the proportion of good degrees has increased as the year of entry increases (hypothesis C) for both male and female students. Figure 9 shows that grade inflation does exist for the proportion of good degrees, providing further evidence in support of hypothesis C.

UCAS entry tariff and interaction with gender. We then asked whether institutions requiring high entry qualifications award a higher proportion of good degrees. Figure 10 shows that this is the case, similar to the proportion of first class degrees.

As was the case with the proportion of first class degrees, the interaction between UCAS tariff and gender is significant for the proportion of good degrees ($p < 0.01$). However, as Figure 10 shows, the magnitude of the effect is small (male slope is 21% steeper than female slope, $p < 0.05$), indicating that females with high entry qualifications earn good degrees with a slightly lower rate than male students. Therefore, this result supports hypothesis A (females perform worse). However, when analyzing upper




Fig. 9. Influence of graduation year on the proportion of good degrees, by gender

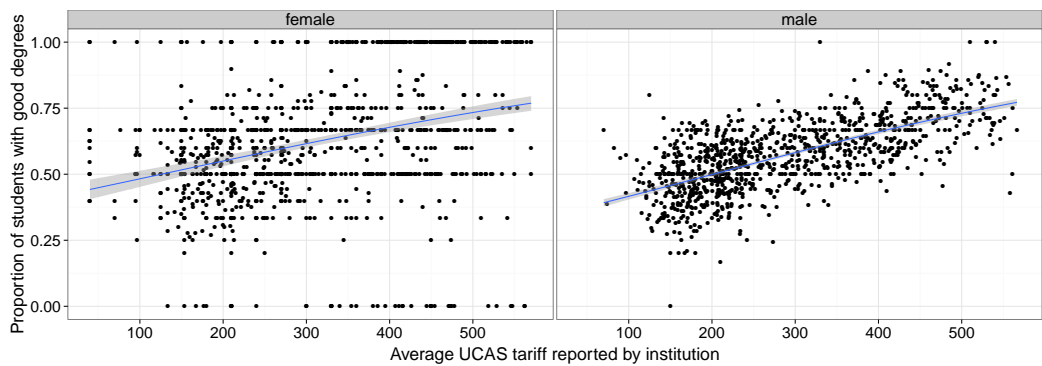


Fig. 10. Influence of UCAS tariff on the proportion of good degrees, by gender

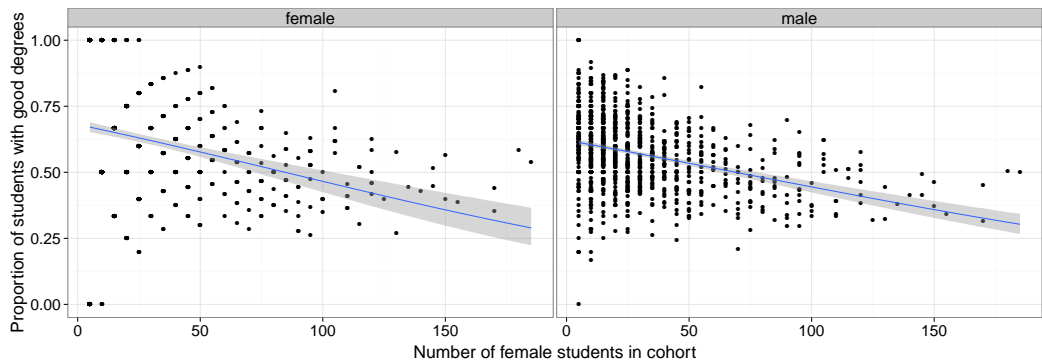


Fig. 11. Influence of the number of females on the proportion of good degrees

second class degrees separately, the interaction between UCAS tariff and gender is not significant ($p < 0.13$), which provides evidence against hypothesis A.

Number of female students. As shown in the case of first class degrees, the data indicates that performance of both male and female students decreases when the number of female students in a program increases (Figure 11). This provides further ev-




Fig. 12. Correlation between the total number of students and the proportion of good degrees




Fig. 13. Correlation between the percentage of female students and the proportion of good degrees




Fig. 14. Proportion of first class degrees in other subjects, by gender

idence against hypothesis B (females perform better when their number/percentage increases).

Other variables. As was the case with the proportion of first class degrees, two explanatory variables were not statistically significant for the proportion of good degrees. The total number of students has a slightly positive, but not significant ($p > 0.22$) correlation with the proportion of good degrees, as shown in Figure 12. Figure 13 shows that the percentage of female students is negatively correlated with the proportion of good degrees, but the result is not statistically significant ($p > 0.58$). This provides further evidence against hypothesis B (females perform better when their number/percentage increases).

In summary, when using the proportion of good degrees (first and upper second class degrees combined) as performance indicator, we have found evidence to support hypothesis C (grade inflation), evidence against hypothesis B (females perform better when their number/percentage increases), and evidence both for and against hypothesis A (females perform worse).

4.3. Contrasting Computer Science with Other Subjects

To check whether the effect observed in CS – a performance gap in the proportion of first class degrees, but not in the proportion of good degrees – also exist in subjects with different gender ratios, we evaluated data from four other subjects. We selected four representative areas which have progressively decreasing proportions of female

students: Subjects Allied to Medicine (a female-dominated subject, 79.95% females in 2013); Business & Administrative Studies (50.44% females in 2013); Mathematical Sciences (40.2% females in 2013); and Engineering & Technology (17.4% females in 2013).

In all four subjects, we found that a positive correlation existed between the graduation year and both the proportion of first class and good degrees (grade inflation); UCAS tariff was positively correlated with the performance indicators (both proportion of first class and good degrees); and females were awarded a higher proportion of good degrees than males. These results are similar to our analysis for CS.

The results for the correlation between gender and the proportion of first class degrees differ significantly from those that we obtained for CS. Figure 14 shows that in all four subjects, females were awarded a significantly higher proportion of first class degrees than males (all $p < 0.001$). Therefore, we found no gender gap in the proportion of first class degrees for Subjects Allied to Medicine, Business & Administrative Studies, Mathematical Sciences, and Engineering & Technology.

Table I. Support for hypotheses A–C, depending on performance indicator

	First class degree	Good degree
A: Female computer science students perform worse than male students	yes	mixed
B: Female students perform better as their total number and percentage increases	no	no
C: Both male and female students perform better as the year of entry increases	yes	yes

5. IMPLICATIONS

The results of this statistical study, summarized in Table I, show that there is a performance gap in the proportion of first class degrees awarded to CS students, but not in the proportion of ‘good’ degrees.

This performance gap may widen when the number of females in computer science increases. According to the cream of the crop effect [Milyo and Schosberg 2000], the low percentage of females results from gender barriers that restrict entry into computer science, such as stereotypes in a society [Cheryan et al. 2015] and reduced exposure to computer science education in high school [Google 2015]. Because the best females join computer science first, an increase in the number of females will reduce average performance and increase the performance gap.

As a step towards addressing the gender and performance gaps, we will first review published explanations for these gaps, followed by an overview of possible solutions.

5.1. Explanations

According to Trauth et al. [2004], there are three main theoretical models to explain the underrepresentation of females in CS. The *essentialist* model assumes that there are inherent, biological differences between males and females which cannot be removed. Thus, the essentialist model posits that men and women should be treated differently on a “separate but equal” basis. The *social construction* model assumes that the gender imbalance is caused by societal forces that shape CS as a male domain. Thus, this model suggests to help women fit into this domain, or to reshape the domain itself. Finally, the *individual differences* model argues that the differences between individuals may be greater than differences between the genders. This model explains the gender gap by a combination of socio-cultural influences, suggesting to address the gender gap by improving each of the influences individually.

In a large meta-study, Ceci et al. [2009] survey biological and socio-cultural explanations for the underrepresentation of females. They conclude that evidence supports a

range of different socio-cultural influences, while the evidence for biological differences is inconclusive.

Many socio-cultural influences have been discussed in the literature, but mainly focusing on the underrepresentation of females. In contrast, here we focus on the performance gap, i.e. the issue that females are awarded fewer first class degrees than males. Therefore, we focus on those socio-cultural influences that have a direct impact on student performance: confidence, stereotypes and bias, and curriculum and context.

We intentionally do not discuss the lack of academic exposure at high school level. Although the lack of exposure contributes to the underrepresentation of women [Carter 2006; Google 2014], it is not a predictor for success [Margolis and Fisher 2002], and its relation to student performance is unclear.

5.1.1. Confidence. Beyer et al. [2003] report that female students have significantly lower confidence in their computer aptitude than males, even when maths scores are taken into account. Fisher et al. [1997] find that female students rate themselves significantly lower in preparedness for their CS courses and ability to master the course material, which leads them to question whether they belong in a CS program. According to Fisher and Margolis [2002], this low confidence is not caused by academic performance, but by unfavorable comparisons with others. All three of the above studies suggest that low confidence decreases the likelihood that women will choose to study CS, and increases the likelihood that they will drop out. Barker et al. [2002] find that CS classrooms exhibit a defensive climate, for example through an impersonal environment and communications that emphasize superiority. This climate is characterized by competitiveness rather than cooperation and can thus impede learning, particularly for female students.

5.1.2. Stereotypes and Bias. Computer science as a subject is stereotypically seen to be male-dominated, with males performing better than females, and successful females being perceived as exceptional (thus leaving the stereotype intact) [Beyer et al. 2003]. In addition, Blickenstaff [2005] argues that CS presents a masculine world view which tends to exclude females. Another stereotype about CS is an expectation of brilliance. Leslie et al. [2015] report that computer scientists believe that some “innate talent” is required for success in their subject. Since women are stereotyped as not having such talent, they are underrepresented. These stereotypes are pervasive in society, and start acting on children early on, for example by inducing pressure on children to fulfill gender roles [Blickenstaff 2005].

Computer scientists are stereotyped as being intelligent, single-mindedly focused on computers, and socially inept – the so-called “geek mythology” [Margolis and Fisher 2002; Beyer et al. 2003]. This stereotype directly conflicts with female gender roles, especially with women’s stronger interpersonal orientation [Beyer et al. 2003], and thus deters women from entering the field of CS [Fisher et al. 1997].

Unconscious bias affects women who have already entered the field, by biasing the decisions about women made by those who assess or hire them. For example, Moss-Racusin et al. [2012] found that both male and female faculty members rate male applicants as significantly more competent and hireable than females, even though application materials were identical.

5.1.3. Curriculum and Context. At the high school level, science curricula tend to focus more on the contributions of male than female scientists [Blickenstaff 2005], which further reinforces the stereotype that women have no place in science. At the university level, Carter [2006] finds that the top reason why women choose CS is their intention to apply computing in another discipline. Margolis and Fisher [2002] also emphasize

the importance of connections to other subjects. However, most CS programs in the UK do not reflect this and rarely focus on applications outside of the core subject.

5.2. Solutions

Many solutions have been discussed to eliminate the gender gap, directed either towards attracting women into CS, or towards retaining them. However, not many solutions are directly applicable to the performance gap. We discuss possible solutions, and how they can contribute to closing the performance gap by raising the proportion of females who are awarded first class degrees.

5.2.1. Curriculum and Course Design. Because females choose CS for its application areas in other disciplines [Carter 2006], designing female-friendly curricula can attract more women to CS. Beyond increasing the number of female students, courses with a broad focus can also improve performance. For example, Rich et al. [2004] report reduced failure rates, Newhall et al. [2014] find that student performance improved by half a letter grade when a broad focus was combined with mentor-led help sessions, and Dodds et al. [2008] find no significant difference in overall course grades between males and females.

A curriculum can be made ‘female-friendly’ by emphasizing the human dimension of computing and its relevance to real life; by allowing creative freedom [Vilner and Zur 2006]; and by drawing examples from a wide range of application areas [Varma 2002]. This can be realized by adding interdisciplinary courses, for example on the design of human-computer interaction [Margolis and Fisher 2002], or by engaging students with non-profit or community organizations, thus providing alternative paths through a CS degree [Fisher and Margolis 2002]. Promising results have been reported with introductory courses that have a broad focus on application areas, for example computational problem solving with problems from science and engineering [Klawe 2013; Alvarado and Dodds 2010]; Internet programming [Kurkovsky 2007]; or media computation [Rich et al. 2004].

Introductory courses with a broad focus are particularly easy to implement in the US, where the major is typically declared after the first year at university, giving students some time to experiment. They are harder to implement in countries where students choose their subject before coming to university. This disadvantage can be offset by adapting marketing materials, asking STEM ambassadors to present a broad interdisciplinary focus when interacting with high school students, and by introducing courses with a broad focus at high school level (the new AP Computer Science Principles course in the US is an example).

5.2.2. Pedagogy. Good teaching can ease entry into a field [Lagesen 2007], but can also ensure that good students are pushed towards top performance. Specifically, females have been shown to benefit from problem-based learning [Lewis et al. 2006], peer support and interaction among classmates [Cohoon 2002], and group work [Vilner and Zur 2006]. Pair programming is a form of group work where two students collaborate on a single computer. The students take turns at the keyboard, while the ‘other’ student spots errors and offers suggestions for alternative design approaches. This has been shown to increase pass rates and help weaker students [Lau and Yuen 2009], and to help female students by increasing their programming confidence [Werner et al. 2004].

5.2.3. Increasing Confidence. Four strategies have been shown to increase the confidence of female students. First, involving female students in research and teaching as early as possible, for example by encouraging them to apply for internships or teaching assistantships [Beyer et al. 2003], involving them in research projects [Cohoon 2002], or offering summer research opportunities after the first year [Alvarado and Dodds

2010]. Second, providing female role models [Vilner and Zur 2006] and mentors [Cohoon 2002; Newhall et al. 2014]. Third, creating a supportive community, for example by sending female students to the Grace Hopper Celebration of Women in Computing [Alvarado and Dodds 2010], creating a web portal [Boschetto and Cortesi 2009], or setting up informal events to strengthen ties between female students and faculty members [Varma 2002]. Fourth, communicating in a positive way, for example by providing information that dispels myths about innate ability [Fisher and Margolis 2002]; publicly expressing positive opinions of female students' strengths [Cohoon 2002]; emphasizing that there is no need to measure themselves against other students [Varma 2002]; or highlighting that sustained effort, not innate giftedness, is most important for success in CS [Leslie et al. 2015].

In addition to confidence, a student's mindset can influence performance. According to Dweck [2000], students who believe that intelligence can be increased through learning (growth mindset) perform better than students who believe that intelligence is fixed (fixed mindset). A growth mindset can be fostered by praising students for effort and persistence instead of smartness [Dweck 2007]. However, in a study by Lewis [2007], 77% of faculty disagreed with the idea of a growth mindset. Making faculty aware of Dweck's research can increase this number and thus improve student performance.

5.2.4. Countering Stereotypes and Biases. Raising awareness of stereotypes and biases is an important step towards closing the performance gap, because faculty members may unconsciously disadvantage female students if they are not aware of stereotypes and unconscious bias [Moss-Racusin et al. 2012].

However, raising awareness is not enough to remove the effect of gender stereotypes and biases when they operate on an unconscious level [Hill et al. 2010]. We highlight two techniques that have been shown to reduce gender stereotypes and biases. First, implementation intentions are action plans that involve thinking counterstereotypical thoughts in situations where hiring or assessment decisions about females are made. This has been shown to reduce automatic stereotype bias [Stewart and Payne 2008]. Implementation intentions can be formulated as if-then statements, for example "when I see a female student, then I will think *competent*." At an individual level, this can reduce unconscious bias in grading female students' assignments.

Second, visual priming can reduce automatic prejudice for at least 24 hours after priming [Dasgupta and Greenwald 2001]. For example, viewing pictures of admired females can reduce automatic prejudice against females. Nonstereotypic association training turns the principle of visual priming into a systematic training program where participants repeatedly assign nonstereotypic traits to male and female faces. This has been shown to reduce stereotype activation [Kawakami et al. 2005]. At a departmental level, visual priming can be achieved by displaying murals or posters of admired female computer scientists [Handelsman and Moss-Racusin 2013].

6. SUMMARY

We have presented a large-scale statistical analysis of student performance in computer science across the UK. We found that a performance gap exists in the proportion of first class degrees, i.e. females were awarded significantly fewer first class degrees than males. We did not find a performance gap for the proportion of 'good' degrees (first and upper second class combined). We investigated four other subjects – Subjects Allied to Medicine, Business & Administrative Studies, Mathematical Sciences, and Engineering & Technology – and found that the performance gap does not exist in these subjects. Based on these findings, we surveyed possible explanations for the performance gap, and solutions to eliminate it. Solutions include the design of curricula

and courses; pedagogical approaches; ways to increase the confidence of female students; and ways to counter stereotypes and unconscious biases. Most solutions require raising awareness among staff, and can be implemented without major institutional changes.

ACKNOWLEDGMENTS

The author thanks the anonymous reviewers whose valuable comments greatly improved the paper.

REFERENCES

- Alvarado, C. and Dodds, Z. (2010). Women in CS: An Evaluation of Three Promising Practices. In *Proceedings of the 41st ACM Technical Symposium on Computer Science Education*, SIGCSE '10, pages 57–61, Milwaukee, WI, USA. ACM.
- Ashcraft, C. and Blithe, S. (2009). Women in IT: The facts. *National Center for Women & Information Technology*.
- Barker, L. J., Garvin-Doxas, K., and Jackson, M. (2002). Defensive Climate in the Computer Science Classroom. In *Proceedings of the 33rd SIGCSE Technical Symposium on Computer Science Education*, SIGCSE '02, pages 43–47, New York, NY, USA. ACM.
- Barker, L. J., McDowell, C., and Kalahar, K. (2009). Exploring Factors That Influence Computer Science Introductory Course Students to Persist in the Major. In *Proceedings of the 40th ACM Technical Symposium on Computer Science Education*, SIGCSE '09, pages 153–157, Chattanooga, TN, USA. ACM.
- Beyer, S., Rynes, K., Perrault, J., Hay, K., and Haller, S. (2003). Gender Differences in Computer Science Students. In *Proceedings of the 34th SIGCSE Technical Symposium on Computer Science Education*, SIGCSE '03, pages 49–53, Reno, NV, USA. ACM.
- Blickenstaff, J. C. (2005). Women and science careers: leaky pipeline or gender filter? *Gender and Education*, 17(4):369–386.
- Boschetto, E. and Cortesi, A. (2009). Women and Informatics: the Ada Web Portal. *Innovation in Teaching and Learning in Information and Computer Sciences*, 8(2):64–72.
- Byrne, P. and Lyons, G. (2001). The Effect of Student Attributes on Success in Programming. In *Proceedings of the 6th Annual Conference on Innovation and Technology in Computer Science Education*, ITiCSE '01, pages 49–52, Canterbury, Kent, United Kingdom. ACM.
- Carter, L. (2006). Why Students with an Apparent Aptitude for Computer Science Don't Choose to Major in Computer Science. In *Proceedings of the 37th SIGCSE Technical Symposium on Computer Science Education*, SIGCSE '06, pages 27–31, Houston, TX, USA. ACM.
- Ceci, S. J., Williams, W. M., and Barnett, S. M. (2009). Women's underrepresentation in science: Sociocultural and biological considerations. *Psychological Bulletin*, 135(2):218–261.
- Cheryan, S., Master, A., and Meltzoff, A. N. (2015). Cultural stereotypes as gatekeepers: increasing girls' interest in computer science and engineering by diversifying stereotypes. *Developmental Psychology*, 6:49.
- Cohoon, J. M. (2002). Recruiting and Retaining Women in Undergraduate Computing Majors. *SIGCSE Bull.*, 34(2):48–52.
- Crawley, M. J. (2012). *The R Book*. Wiley-Blackwell, Chichester, West Sussex, United Kingdom, 2nd edition.
- Dasgupta, N. and Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals.

- Journal of Personality and Social Psychology*, 81(5):800–814.
- Dodds, Z., Libeskind-Hadas, R., Alvarado, C., and Kuenning, G. (2008). Evaluating a Breadth-first Cs 1 for Scientists. In *Proceedings of the 39th SIGCSE Technical Symposium on Computer Science Education*, SIGCSE '08, pages 266–270, Portland, Oregon, USA. ACM.
- Dweck, C. (2007). *Mindset: The New Psychology of Success*. Ballantine Books, New York.
- Dweck, C. S. (2000). *Self-theories: Their Role in Motivation, Personality, and Development*. Psychology Press.
- Fisher, A. and Margolis, J. (2002). Unlocking the Clubhouse: The Carnegie Mellon Experience. *SIGCSE Bull.*, 34(2):79–83.
- Fisher, A., Margolis, J., and Miller, F. (1997). Undergraduate Women in Computer Science: Experience, Motivation and Culture. In *Proceedings of the 28th SIGCSE Technical Symposium on Computer Science Education*, SIGCSE '97, pages 106–110, San Jose, CA, USA. ACM.
- Gharibyan, H. and Gunsaulus, S. (2006). Gender Gap in Computer Science Does Not Exist in One Former Soviet Republic: Results of a Study. In *Proceedings of the 11th Annual SIGCSE Conference on Innovation and Technology in Computer Science Education*, ITICSE '06, pages 222–226, Bologna, Italy. ACM.
- Gibbs, G. (2010). *Dimensions of quality*. Higher Education Academy York.
- Google (2014). Women Who Choose Computer Science—What Really Matters. Technical report. https://docs.google.com/a/google.com/file/d/0B-E2rcvhn1Q_a1Q4VUxWQ2dtTHM/edit?pli=1&usp=embed_facebook. Accessed: 2015-11-20.
- Google (2015). Images of Computer Science: Perceptions Among Students, Parents and Educators in the U.S. Technical report. <https://services.google.com/fh/files/misc/images-of-computer-science-report.pdf>. Accessed: 2015-11-20.
- Handelsman, J. and Moss-Racusin, C. (2013). Laboratory life: Scientists of the world speak up for equality. *Nature*, 495(7439):38.
- Hill, C., Corbett, C., and St. Rose, A. (2010). *Why So Few? Women in Science, Technology, Engineering, and Mathematics*. American Association of University Women.
- Ivanović, M., Putnik, Z., Šišarica, A., and Budimac, Z. (2010). A Note on Performance and Satisfaction of Female Students Studying Computer Science. *Innovation in Teaching and Learning in Information and Computer Sciences*, 9(1):32–41.
- Kawakami, K., Dovidio, J. F., and van Kamp, S. (2005). Kicking the habit: Effects of nonstereotypic association training and correction processes on hiring decisions. *Journal of Experimental Social Psychology*, 41(1):68–75.
- Klawe, M. (2013). Increasing Female Participation in Computing: The Harvey Mudd College Story. *Computer*, 46(3):56–58.
- Konvalina, J., Wileman, S. A., and Stephens, L. J. (1983). Math Proficiency: A Key to Success for Computer Science Students. *Commun. ACM*, 26(5):377–382.
- Kurkovsky, S. (2007). Making Computing Attractive for Non-majors: A Course Design. *J. Comput. Sci. Coll.*, 22(3):90–97.
- Lagesen, V. A. (2007). The Strength of Numbers Strategies to Include Women into Computer Science. *Social Studies of Science*, 37(1):67–92.
- Lau, W. W. F. and Yuen, A. H. K. (2009). Exploring the effects of gender and learning styles on computer programming performance: implications for programming pedagogy. *British Journal of Educational Technology*, 40(4):696–712.
- Leslie, S.-J., Cimpian, A., Meyer, M., and Freeland, E. (2015). Expectations of brilliance underlie gender distributions across academic disciplines. *Science*, 347(6219):262–265.

- Lewis, C. (2007). Attitudes and Beliefs About Computer Science Among Students and Faculty. *SIGCSE Bull.*, 39(2):37–41.
- Lewis, S., McKay, J., and Lang, C. (2006). The Next Wave of Gender Projects in IT Curriculum Teaching at Universities. In *Proceedings of the 8th Australasian Conference on Computing Education - Volume 52*, ACE '06, pages 135–142, Darlinghurst, Australia. Australian Computer Society, Inc.
- Margolis, J. and Fisher, A. (2002). *Unlocking the clubhouse: Women in computing*. MIT press.
- Margolis, J., Fisher, A., and Miller, F. (2000). The Anatomy of Interest: Women in Undergraduate Computer Science. *Women's Studies Quarterly*, 28(1/2):104–127.
- McGill, R., Tukey, J. W., and Larsen, W. A. (1978). Variations of Box Plots. *The American Statistician*, 32(1):12–16.
- McKenzie, K. and Schweitzer, R. (2001). Who Succeeds at University? Factors predicting academic performance in first year Australian university students. *Higher Education Research & Development*, 20(1):21–33.
- Mellström, U. (2009). The Intersection of Gender, Race and Cultural Boundaries, or Why is Computer Science in Malaysia Dominated by Women? *Social Studies of Science*, 39(6):885–907.
- Milyo, J. and Schosberg, S. (2000). Gender Bias and Selection Bias in House Elections. *Public Choice*, 105(1-2):41–59.
- Moss-Racusin, C. A., Dovidio, J. F., Brescoll, V. L., Graham, M. J., and Handelsman, J. (2012). Science faculty's subtle gender biases favor male students. *Proceedings of the National Academy of Sciences*, 109(41):16474–16479.
- Newhall, T., Meeden, L., Danner, A., Soni, A., Ruiz, F., and Wicentowski, R. (2014). A Support Program for Introductory CS Courses That Improves Student Performance and Retains Students from Underrepresented Groups. In *Proceedings of the 45th ACM Technical Symposium on Computer Science Education, SIGCSE '14*, pages 433–438, Atlanta, GA, USA. ACM.
- Papastergiou, M. (2008). Are Computer Science and Information Technology still masculine fields? High school students' perceptions and career choices. *Computers & Education*, 51(2):594–608.
- Prey, J. C. and Weaver, A. C. (2013). Fostering Gender Diversity in Computing. *Computer*, 46(3):22–23.
- Rich, L., Perry, H., and Guzdial, M. (2004). A CS1 Course Designed to Address Interests of Women. In *Proceedings of the 35th SIGCSE Technical Symposium on Computer Science Education, SIGCSE '04*, pages 190–194, Norfolk, VA, USA. ACM.
- Shillabeer, A. and Jackson, K. (2013). Gender Imbalance in Undergraduate IT Programs – A Vietnamese Perspective. *Innovation in Teaching and Learning in Information and Computer Sciences*, 12(1):70–83.
- Stewart, B. D. and Payne, B. K. (2008). Bringing Automatic Stereotyping Under Control: Implementation Intentions as Efficient Means of Thought Control. *Personality and Social Psychology Bulletin*, 34(10):1332–1345.
- Trauth, E. M., Quesenberry, J. L., and Morgan, A. J. (2004). Understanding the Under Representation of Women in IT: Toward a Theory of Individual Differences. In *Proceedings of the 2004 SIGMIS Conference on Computer Personnel Research: Careers, Culture, and Ethics in a Networked Environment, SIGMIS CPR '04*, pages 114–119, Tucson, AZ, USA. ACM.
- Varma, R. (2002). Women in Information Technology: A Case Study of Undergraduate Students in a Minority-Serving Institution. *Bulletin of Science, Technology & Society*, 22(4):274–282.
- Vilner, T. and Zur, E. (2006). Once She Makes It, She is There: Gender Differences in Computer Science Study. In *Proceedings of the 11th Annual SIGCSE Conference*

- on Innovation and Technology in Computer Science Education*, ITICSE '06, pages 227–231, Bologna, Italy. ACM.
- Werner, L. L., Hanks, B., and McDowell, C. (2004). Pair-programming Helps Female Computer Science Students. *J. Educ. Resour. Comput.*, 4(1).
- Wilson, B. C. (2002). A Study of Factors Promoting Success in Computer Science Including Gender Differences. *Computer Science Education*, 12(1-2):141–164.
- Yorke, M. (2009). *Trends in honours degree classifications, 1994-95 to 2006-07, for England, Wales and Northern Ireland*. The Higher Education Academy, York.