

Evaluating Maps in a Massive Open Online Course

Anthony C. Robinson
The Pennsylvania State University
arobinson@psu.edu

Jonathan K. Nelson
The Pennsylvania State University
jkn128@psu.edu

New forms of cartographic education are becoming possible with the synthesis of easy to use web GIS tools and learning platforms that support online education at a massive scale. The internet classroom can now support tens of thousands of learners at a time, and while some common types of assessments scale very easily, others face significant hurdles. A particular concern for the cartographic educator is the extent to which original map designs can be evaluated in a massive open online course (MOOC). Based on our experiences in teaching one of the first MOOCs on cartography, we explore the ways in which very large collections of original map designs can be assessed. Our methods include analysis of peer grades and qualitative feedback, visual techniques to explore design methods, and quantitative comparison between expert ratings and peer grades. The results of our work suggest key challenges for teaching cartography at a scale where instructors cannot provide individual feedback for every student.

KEYWORDS: cartographic education; MOOCs; online learning; peer assessment

INTRODUCTION

A NEW SPIRIT of institutional openness, coincident with the emergence of new forms of education via the internet, has combined to drive the development of learning experiences that reach massive, global audiences. The massive open online course (MOOC) is one such example, growing from an initial pedagogical experiment with two thousand students in 2008 (McAuley et al. 2010) to mature platforms today featuring hundreds of courses from universities around the world for an audience measured in the tens of millions (Pappano 2012). At the same time, mapping technology has proliferated to reach enormous new audiences through location-enabled mobile devices and easy to use web mapping tools. As a result, cartographers have the unique opportunity today to reach massive, global audiences through learning experiences at scale.

The potential to teach cartography to thousands, rather than dozens, has immediate attraction to cartographic educators with an eye on encouraging a broader public understanding of best practices in map reading and design. It also introduces significant new challenges to overcome. We explore one of those challenges here by evaluating the extent to which map design assessment can take place in a massive, distributed global classroom. If we intend to expand the range of students who engage with cartography

through increased openness, then we must address the fundamental issue of scale between the relatively few capable cartography educators in the world, compared to the very large potential audience of mapmakers who may be keen to learn.

In the sections to come, we begin by characterizing the state of the art in online teaching at scale. As part of this discussion, we focus on previous attempts to teach massive courses within the discipline of geography. Next we describe the use of peer assessment methods, which are the most common means for supporting evaluation of student-generated projects in massive courses.

Using evidence gathered from teaching a MOOC on cartography, we follow our literature review with a methodological structure we have used to explore the reliability and utility of peer assessment through quantitative and visual analysis. The results of these analyses are then discussed and situated within the broader context of challenges and opportunities for scaling cartographic education to massive audiences. We conclude with ideas for future research to explore emerging dimensions of assessment in a new realm of massive online cartographic education.



© by the author(s). This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

THIS ARTICLE IS WRITTEN at a time in which distance education methods and mapping technology have become blendable and distributable in radical new ways. This potential has been a long time in the making, however, through decades of previous development in both areas. The rise of e-learning has roots reaching as far back as the 1960s with early experiments fusing computing with education (Nicholson 2007). The science and technology of e-learning saw its renaissance, however, during the 1990s as delivery via the internet became possible for an increasingly large audience of learners.

Distance education today can take many forms, including fully-online and blended models of instruction which employ synchronous as well as asynchronous types of engagement through assignments, discussion, and content delivery (Unwin et al. 2011). The science of online instruction has also seen major developments and has been the subject of significant attention within geography itself (Clark, Monk, and Yool 2007; Terry and Poole 2012). Evidence from hundreds of controlled studies has helped reveal that design guidelines for effective online learning can be developed, and that courses designed with those imperatives in mind perform as well as their in-person counterparts. Furthermore, online classes can offer unique advantages to students in terms of flexibility and access to match a broad range of potential learning styles (Means et al. 2010).

Fully-online geography courses focused on the mapping and geographic information sciences began to appear at universities and colleges in the late 1990s, starting a trend which continues today to emphasize geospatial technology through online certificate and degree programs (McMaster et al. 2007). Classes offered in these programs may use a range of instructional models, including synchronous and asynchronous content delivery, discussion systems, lecture videos, and virtual laboratories (Unwin et al. 2011). While online learning initially seemed to promise the ability to lower costs and support larger cohorts in classes, these myths have largely been debunked by researchers (DiBiase and Rademacher 2005). Instead, the primary advantages of online learning today have to do with access, as there are millions of learners around the world who cannot attend in-person classrooms. This is particularly an issue for adult education, an area where online programs have grown very rapidly in the United States. Few professionals can relocate to attend on-campus

experiences, and even when one is located nearby, not many can attend evening and weekend courses for weeks on end without interruption.

In the late 2000s, the first experiments began with a new approach to tackle the scale issue associated with online learning. George Siemens and Stephen Downes at the University of Manitoba launched a new online course in 2008 titled *Connectivism and Connective Knowledge*, which was opened for anyone in the world to participate in for free via the internet. This experimental approach drew in more than two thousand learners from all over the world. *Connectivism and Connective Knowledge* is credited today as the first example of a massive open online course (McAuley et al. 2010). Soon, others began to experiment with developing MOOCs and new platforms for creating and delivering MOOCs. These new platforms included Coursera, launched by Andrew Ng and Daphne Koller from Stanford University, Udacity, created by Sebastian Thrun at Stanford University, and edX, which was co-developed by the Massachusetts Institute of Technology and Harvard University. Today, these platforms and many other new entrants are offering hundreds of classes to millions of learners. Coursera is currently the largest MOOC provider, with more than 10 million students taking courses on its platform by late 2014 (Larson 2014). Generally speaking, MOOC platforms develop partnerships with universities to develop and deliver courses, with the platform providers offering their learning management systems and user base, and the university partners providing content and instruction.

There is no doubt that what encourages learning at a massive scale is the fact that MOOCs are normally free to take. While on the surface offering a free course may seem to benefit neither universities nor the MOOC platform providers, new models for revenue generation are emerging through MOOCs via the provision of microcredentials for a small fee, and via lead generation to encourage MOOC students to enroll in traditional tuition-paid online and residential learning programs. Those who support the evolution of MOOCs have pointed out the potential to reach large and globally diverse audiences that are not typically able to access higher education experiences. Those who are critical of the trend highlight the lack of sustainable revenue generation models, low class retention rates, and pedagogical concerns given the fact that instructors

cannot possibly provide individualized instruction and feedback with thousands of students at once.

Since 2013, several MOOCs on geography topics have been developed and taught across a variety of MOOC platforms. The first of these, a course called *Maps and the Geospatial Revolution* (hereafter also referred to as the Maps MOOC), was launched on the Coursera platform in February 2013 (Robinson et al. 2015). Subsequently, several other MOOCs on geospatial science and technology topics have been developed, including *Introduction to GIS using Quantum GIS* (www.canvas.net/browse/delmarcollege/courses/introduction-to-geospatial-technology-1), *Geodesign: Change Your World* (www.coursera.org/course/geodesign), *From GPS and Google Maps to Spatial Computing* (www.coursera.org/course/spatialcomputing), *Geospatial Intelligence and the Geospatial Revolution* (www.coursera.org/course/geoint), and *Going Places With Spatial Analysis* (www.esri.com/landing-pages/training/spatial-analysis). What these courses have in common is that they are targeting new audiences of geographic learners with free experiences to introduce key geospatial topic areas. *Maps and the Geospatial Revolution* appears so far to be the only MOOC that focuses explicitly on cartographic education, though we anticipate there to be many more options in this area in the near future.

The current state of the art in MOOC platform development supports a limited palette of assessment types. Compared to traditional online learning with small cohorts where individualized grading by an instructor is possible, MOOC assessments are normally limited to autograded quizzes and exams using multiple choice or

true/false questions. The key exception to autograded assessments in MOOCs is through the use of peer evaluation frameworks. Many MOOC platforms provide peer assessment tools as the primary means by which individual projects can be evaluated at scale to provide formative feedback.

Peer assessment employs a simple concept at its core: students evaluate the work of their peers (Falchikov and Goldfinch 2000). In a traditional course, peer assessment is frequently employed as a means to generate peer-to-peer discussion on course content or project deliverables. Peer assessment is usually moderated by the instructor in this setting to ensure that feedback is constructive and consistent. Peer assessment has been widely adopted in MOOCs to overcome the problem that MOOC assessment methods are otherwise limited to summative measures that come from autograded quizzes and exams (Suen 2014). In contrast to its application in non-massive courses, peer assessment is not easily moderated by an instructor in a MOOC where there may be thousands of assignments and reviewers at work.

In the sections that follow, we describe the development of a peer assessment intended to support formative feedback for individual cartographic design projects generated in a MOOC. Using data collected from students completing this assignment, we explore several ways in which the resulting grades and assignment can be evaluated to gauge the challenges and opportunities that this framework poses for further explorations in teaching cartography at a massive scale.

PEER ASSESSMENT IN THE MAPS MOOC

TO EXPLORE THE POTENTIAL for map assessment in courses at scale, we have analyzed map contributions from students enrolled in *Maps and the Geospatial Revolution*. Since 2013, this MOOC has been taught three times by the first author, Anthony Robinson, enrolling more than 100,000 students from over 200 countries and territories.

To evaluate peer assessment reliability and consistency, we use assessments and peer ratings from the first Maps MOOC taught in July, 2013. That session enrolled over 49,000 students, with more than 36,000 participating once the class became active. 3,064 earned a passing grade

in the course, with more than 8,000 students active during its final week. High attrition rates are common across MOOCs (Ho et al. 2014), and yet it would take over a hundred sections of a typical cartography course to reach even the relatively small subset who earned a passing grade in the initial offering of the Maps MOOC.

Students in the 2013 session of the Maps MOOC created a total of 2,787 final project submissions. This project represented the culminating effort for the class, and it required students to find spatial data (or create it themselves) and create their own original maps to tell stories

about their chosen data sets. Three potential options were presented, from more difficult to less difficult, depending on students' self-assessment of their mapping skill levels. Option 1 suggested the use of Esri's ArcGIS Online tools, which students use in four lab assignments in the course, and therefore the easiest option for those who have made maps of their own only as a result of taking the class. Option 2 suggested the use of Esri's StoryMap templates and tools, which requires technical skill beyond what is explicitly taught in the MOOC. Option 3 encouraged students to use CartoDB, MapBox Studio, or a desktop GIS such as QGIS to complete their projects. This option was intended to push the more experienced students taking the class to build upon their existing knowledge.

RUBRIC AND ASSESSMENT DESIGN

A critical pedagogical component for the design of any peer assessment is its grading rubric. Rubric design deserves special attention in a MOOC given the very wide range of backgrounds and expertise evident in their globally-diverse audiences. In addition, we know that MOOCs typically have more than 50% of their student populations speaking a primary language aside from English, so the language used to describe rubric elements must take this into account to the extent it is possible.

In conjunction with a learning designer, who helped advise on the development of the Maps MOOC content and assignments, we developed a four-part rubric that asks

Does this map tell a complete story?

- 0 - this map does not appear to tell any story at all.
- 1 - this map begins to tell a story, but falls short of offering a complete narrative.
- 2 - this map tells a good story, but some aspects would benefit from further clarification.
- 3 - this map tells a complete story.

Does this map tell a compelling story?

- 0 - this map does not tell a compelling story at all.
- 1 - this map tells a story that is partially compelling.
- 2 - this map tells a story that is mostly compelling.
- 3 - this map tells a very compelling story.

Is this map designed in a way that reflects the use of best practices in cartographic design and geospatial analysis?

- 0 - this map shows no attention to best practices in cartographic design and geospatial analysis.
- 1 - this map shows only a little attention to best practices in cartographic design and geospatial analysis.
- 2 - this map shows a solid level of attention to best practices in cartographic design and geospatial analysis.
- 3 - this map follows all best practices in cartographic design and geospatial analysis.

Does this map have an aesthetic look and feel that reinforces the objectives of the story it tells?

- 0 - this map has a look and feel that has no evident relation to the story it tells.
- 1 - this map has a look and feel that has only a little relation to the story it tells.
- 2 - this map has a look and feel that has some relation to the story it tells.
- 3 - this map has a look and feel that relates very well to the story it tells.

Figure 1.

peer graders to evaluate how well each submission presents a complete story, how compelling that story is, whether or not the map design uses best practices in cartographic design, and the extent to which the map has an aesthetic look and feel that reinforces its storytelling objectives. Each of these four elements could be rated from 0 to 3, as shown in the detailed rubric in Figure 1. The maximum possible grade for this assignment is therefore 12 points, and the assignment was weighted to make up 20% of the overall grade for the class.

The Coursera platform allows an instructor to specify how many submissions each student must grade. We chose to require a minimum of three graded submissions for each participant. To our surprise, most students voluntarily graded additional submissions. In fact, a total of 1,825 submissions received at least five peer grades. The Coursera platform allows students to voluntarily grade more than the required number of submissions, and it uses the median rather than the mean of scores to determine the final grade for a peer review assignment.

RESULTS

TO EVALUATE THE EXTENT to which map design can be assessed at scale, in this section we explore multiple aspects of peer grading results from the Maps MOOC. We begin by describing the general outcomes seen across three sessions of teaching the course. Next we provide evidence from a quantitative evaluation of the stability and reliability associated with peer grading from the first session of the Maps MOOC. Finally, we show how techniques and tools from image analysis can be used to begin exploring the qualitative dimensions of map designs submitted in a massive course, using the second Maps MOOC project collection as an example.

PEER GRADING ACROSS ALL SESSIONS

One mechanism for comparing peer assessment results across all three sessions of the Maps MOOC is to proportionally summarize grading for each class. Figure 2 shows peer grade distributions across eleven score ranges for each of the three Maps MOOCs taught so far. The average score for each session is also plotted as a line in a corresponding color on this graph.

The 2013 session featured a broader and flatter distribution of score ranges, particularly between 50% and 90%, compared to the subsequent 2014 and 2015 sessions, which have strong peaks in the 80% to 90% range. Although

the core course content has remained the same across all three sessions, improvements have been made each time to the instructions provided for the peer assessment activity, as it is one element of the course which appears to be the toughest for students to understand and execute compared to the autograded quizzes, lecture videos, and other materials. We suspect that improvements in how the assignment is presented and explained may help students toward developing higher quality submissions that better fit the rubric imperatives. It is also possible that the size of the cohort plays a significant role, as the first MOOC in 2013 had roughly twice as many participants as the 2014 and 2015 sessions.

On average, scores have increased with each subsequent session by a small, but notable margin. Again, while we cannot be certain of the cause, one potential reason for this would be the constant improvements we have made to the assignment instructions to explain deadlines, the peer grading method itself, and how the rubric should be interpreted/employed. Another potential explanation is that MOOC students in general are becoming more familiar and comfortable with peer grading as a common element of massive courses. We note that more students appeared to struggle with this concept in the 2013 class than in subsequent courses, based on what we have seen students discuss in the forums about this assignment.

% of Learners in Each Score Range




Figure 2.

COMPARISON TO INSTRUCTOR GRADING

To evaluate the extent to which peer grading correlates with expert grading by a qualified instructor, we manually graded a 5% random sample (93) from the set of submissions that had received at least five grades ($n=1825$) from the first session of the Maps MOOC taught in 2013. The first author reviewed and graded each submission from this subset using the same rubric as used by the students for peer assessment and was blinded to the grades that had already been assigned by students.

Using this manually-graded set of assignments as a sample, we were able to evaluate the reliability and validity of peer grading. We define *reliability* as the tendency for peer graders to agree with each other when rating a given assignment, and we define *validity* here as the agreement between student-provided grades with expert-provided grades.

Reliability evaluation using intraclass correlation coefficient (ICC) analysis reveals that while agreement among individuals is low ($ICC = 0.262$), taking the averages of five scores provides significant improvements to reliability ($ICC = 0.640$). Our evaluation of score validity using Pearson's correlation coefficient shows that instructor grades have a strong positive correlation with peer grades provided by students ($r = 0.619$, $p > 0.01$). Further details on our analysis of reliability and validity, including the results of a student survey to evaluate the extent to which students understand and appreciate peer grading, can be found in a recent complementary article (Luo, Robinson, and Park 2014).

VISUAL ANALYSIS

While quantitative evaluation provides insights regarding the overall reliability and utility associated with peer grading in cartography courses at scale, this approach completely obscures the artifacts themselves. What should instructors do if they want to actually see and understand the map designs that students have created in such a large course environment? If thousands of maps are created and submitted, how can cartographic educators make sense of what was made beyond basic measures of overall grades and their reliability?

With this motivation in mind, we set out to explore the visual design of submissions from the second session of the

Maps MOOC taught in the spring of 2014. This course generated 1,243 final project submissions from students.

To begin evaluating the look of these maps, first we manually captured screenshots from every map submission and coded them into categories according to the tools used in their creation. Most of the maps (91%) utilized a type of Esri tool (ArcGIS Online, StoryMaps, etc.), while the remaining 9% utilized an alternative mapping platform (CartoDB, Mapbox, Google Maps, etc.). In the context of map design evaluation at scale, this is an important attribute because it highlights the key media used to generate cartographic products across the globe, while it also assesses the extent to which students are applying the tools taught in the MOOC to make maps. This knowledge can help guide future course offerings by suggesting relevant tools to introduce students to. Moreover, the tools used to create maps, their popularity, and their ease of use all have a significant influence on map aesthetics and design processes.

To explore the visual characteristics of these maps, we adopted techniques from image analysis and utilized the ImageJ toolkit (Schneider, Rasband, and Eliceiri 2012), which allows us to combine qualitative categorizations we encode for submissions with automated evaluation of high-level image features extracted from screenshots of the map submissions. Essentially, each map image is represented by four attributes: the tool used to create the map; median saturation value; median brightness value; and median hue value. Given these attributes, we can plot the entire collection of map images in two-dimensional spaces to illuminate visual signatures of map design at both the *global* (entire class) and *local* (individual student) levels.

Figure 3 depicts a montage of the entire collection of student maps, grouped by the different tools used for map creation and sorted darkest to brightest, from left to right, based on median brightness values. This montage conveys the distributions of maps by software type and highlights map types that tend to be brighter or darker overall. At full resolution, one can pan and zoom on the montage to explore and evaluate map designs at the individual level, as a collection, or within/ between mapping software groups.

For example, the most widely used Esri ArcGIS Online maps, shown in the uppermost block of Figure 3, tend to be brighter overall. In contrast, the Esri StoryMap submissions, shown in the fifth block down, are considerably

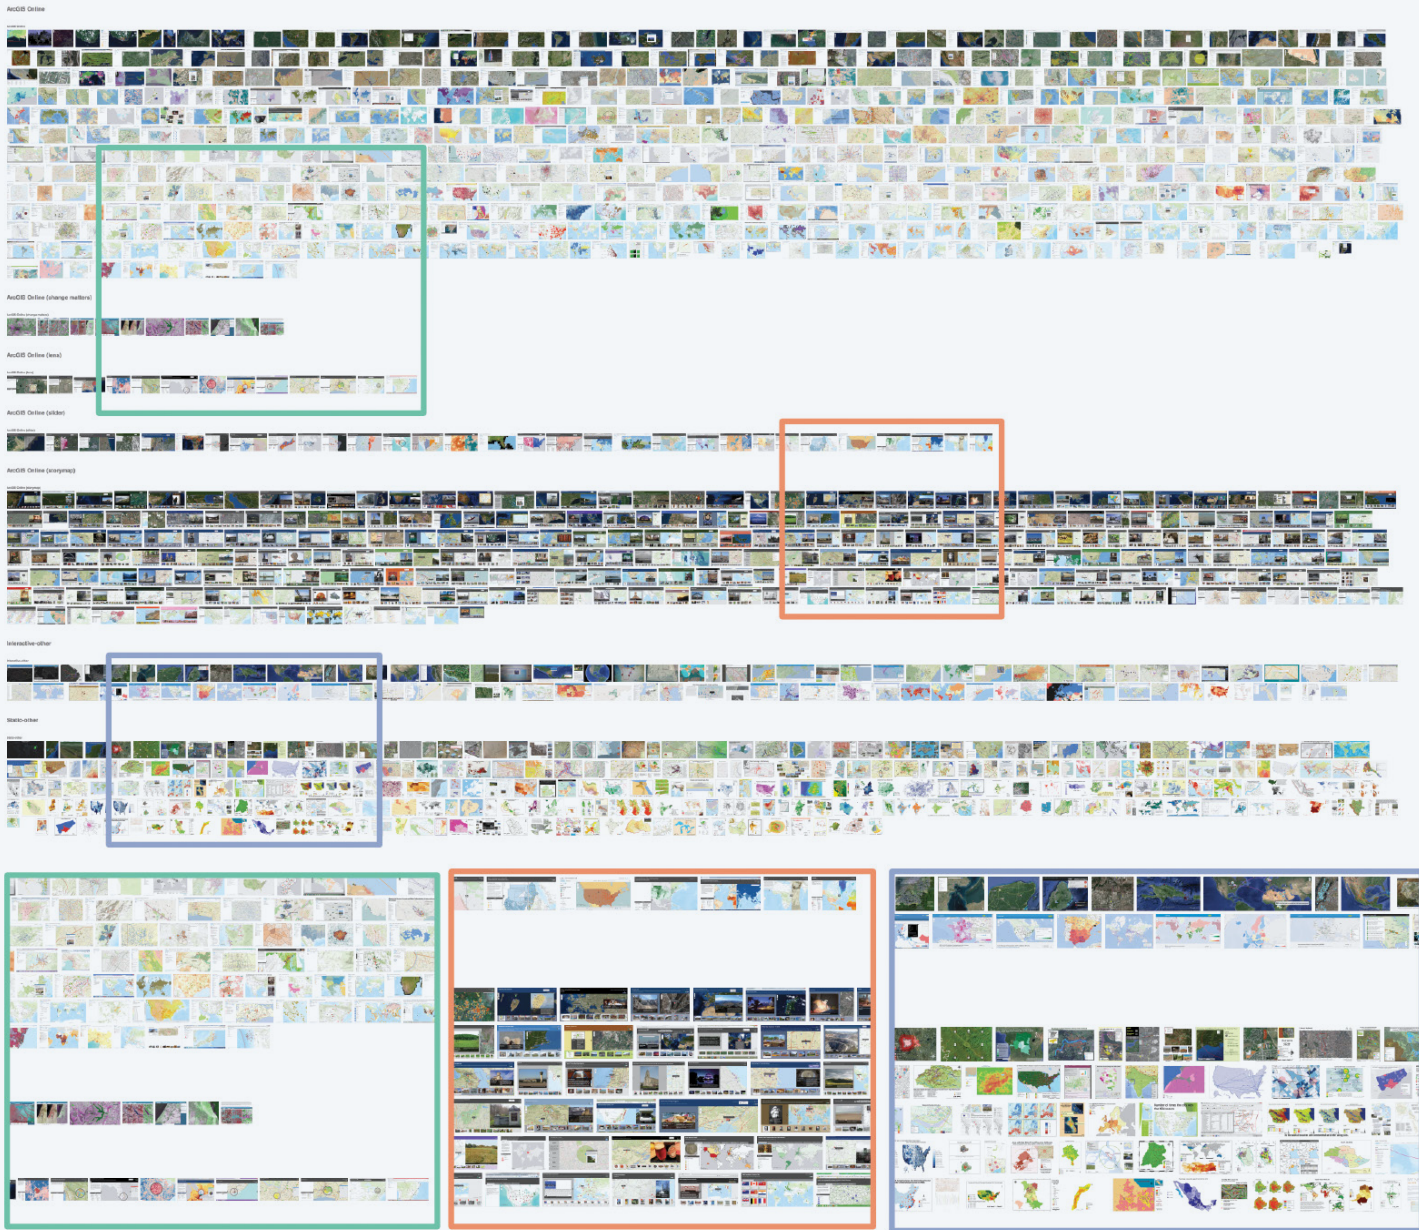


Figure 3.

darker. A closer look at maps in these two categories reveals that students who used Esri ArcGIS Online maps tended to map larger areas, essentially presenting information at smaller map scales, which resulted in more ocean coverage and brighter base map elements. Students who used Esri StoryMaps tended to focus on very specific places, presenting their stories at large map scales and integrating photographs to provide rich context. The lack of ocean and brighter base map elements resulted in overall darker maps. These insights allow cartographic educators to better understand the motivations behind students' individual design choices as well as the role of mapping software in shaping design decisions and overall aesthetics.

Another approach to visualizing map design is to plot map images in a scatterplot using values associated with their visual features. Figure 4 plots map images by median brightness values on the horizontal axis and median saturation values on the vertical axis. The concentration of map images in the bottom right corner of the plot illustrates the strong tendency for students to design bright, unsaturated maps. This trend seems to align with both cartographic theory and with the default map layouts in spatial media authoring software designed with cartographic theory in mind. These maps take visual hierarchy into consideration. The visual characteristics of the base map data which, in most cases, are most influential on the

high-level visual features extracted from the map images are subtle and bright. The darker, more saturated colors are used sparingly in these maps to bring primary data to the top of the visual hierarchy.

At the inverse end of the plot, map images are dark, saturated, and typically representative of designs that use

satellite imagery as a base map. Maps located more centrally in the plot tend to be vector/raster mashups, Esri ArcGIS Online story maps that integrate photographs into the map design, or large-scale maps composed primarily of landmass. Outliers in the scatterplot may represent map designs that are novel, or that could benefit from constructive critique. From an evaluation perspective, the

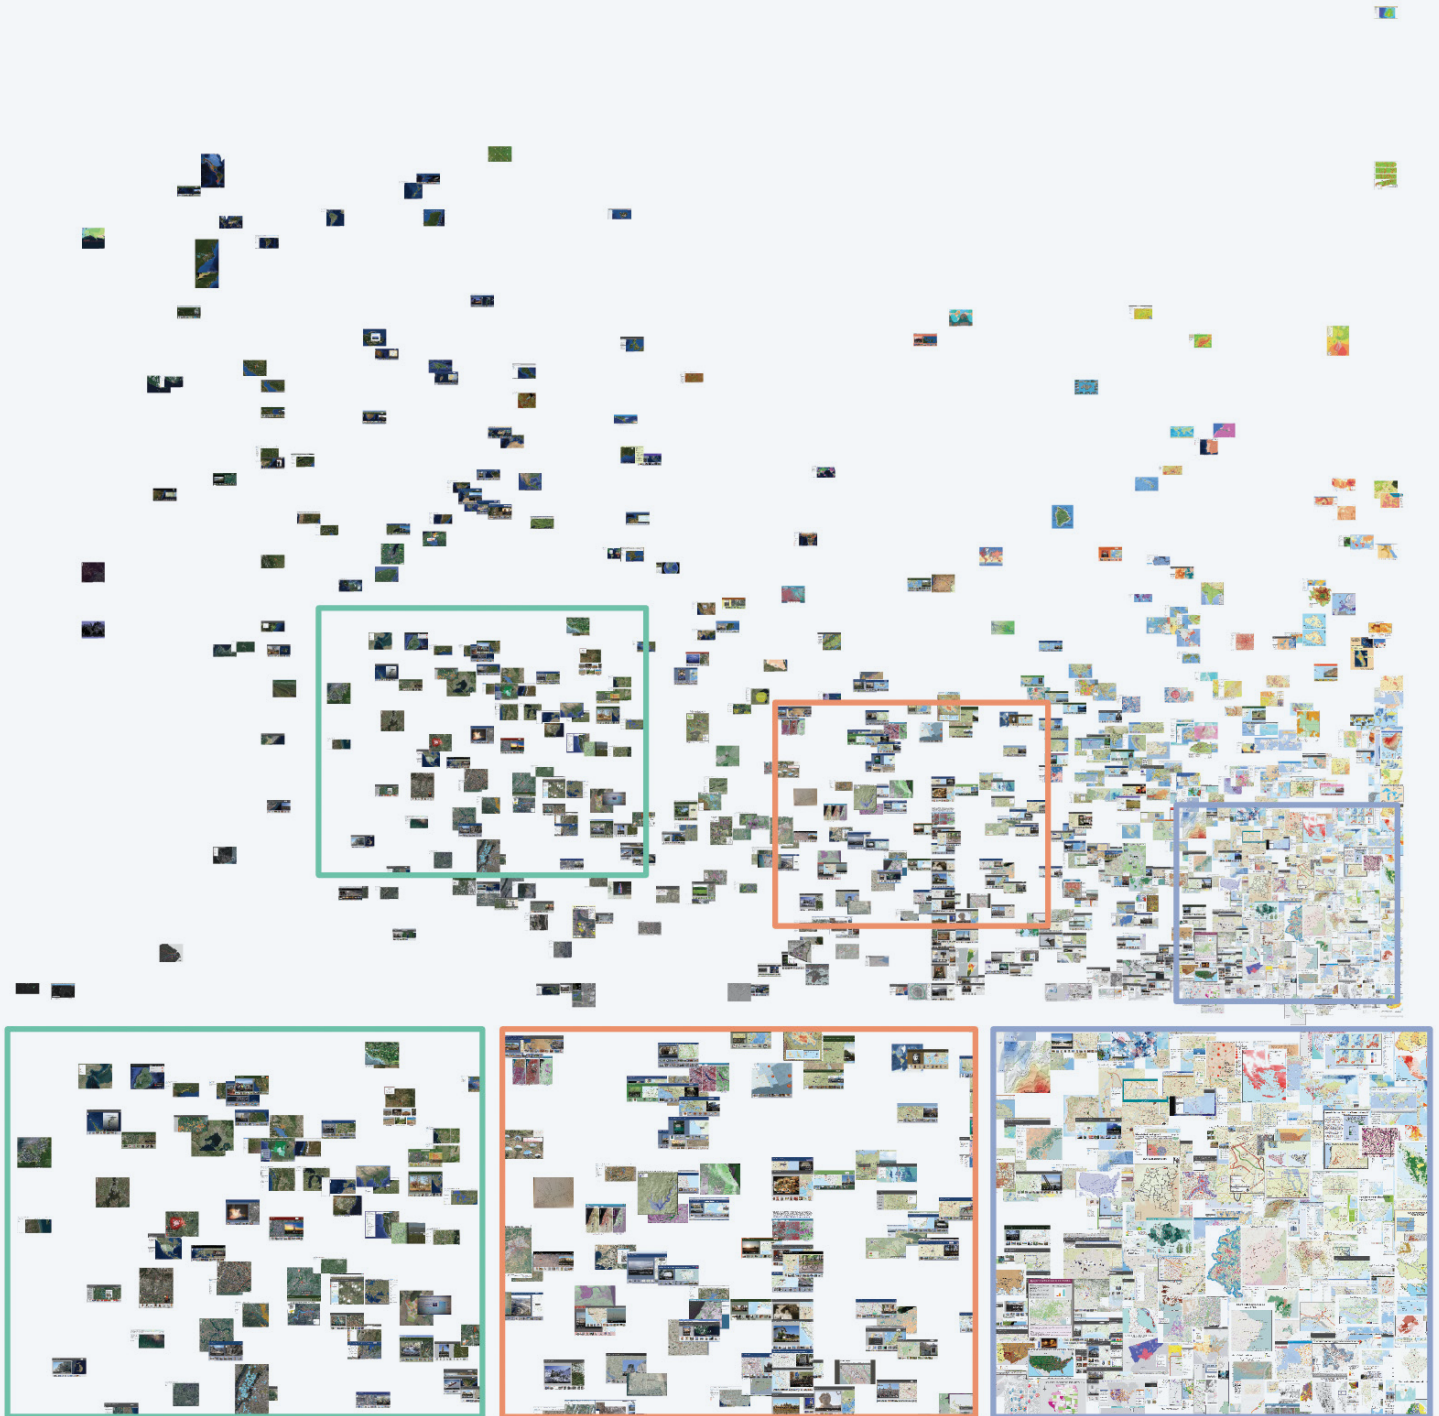


Figure 4.

scatterplot serves as a tool that allows educators to assess students' individual design decisions on visual hierarchy together with software's influence on realizing those

decisions. We explore additional methods for visual analysis of peer-assessed map designs in Nelson and Robinson (2015).

CHALLENGES FOR MAP ASSESSMENT AT SCALE

BASED ON OUR RESULTS from evaluating map designs through quantitative and qualitative means as shown in the previous sections, we propose a series of new research challenges for cartographers to address in order to support map assessment at scale.

WHAT CAN BE DONE TO SUPPORT ITERATIVE MAP DESIGN AND PROGRESSIVE FEEDBACK AT SCALE?

Current peer assessment methods in MOOCs do not support iterative feedback and project development, making it hard to envision a cartographic design course that goes deeper in the way that most cartographic educators would desire. To move beyond single-stage peer assessment in a course at scale would require the development of new platforms that can organize multi-stage reviewing automatically, as well as rubrics that take iteration into account automatically. In a typical cartography class, an instructor will normally assume prior knowledge gains as a class goes from week to week, and penalties for problems may increase over time, while expectations for attention to detail also increase.

While this challenge is a significant one to tackle, it is worth noting that students are already engaging in iterative refinement through informal means in a class like the Maps MOOC. We have observed students posting projects in progress to the discussion forum and soliciting critique for multiple drafts over a period of days or weeks to improve their submissions. This promising sign is tempered by the fact that these students are engaging in ungraded peer review without a standardized rubric. These are two key aspects of peer assessment that would need to be adapted to support iterative progress in a formal assignment.

HOW CAN EDUCATORS SEE AND UNDERSTAND LARGE COLLECTIONS OF MAP SUBMISSIONS?

As we have shown here, it is possible to begin making sense of very large collections of map submissions through

the use of image analysis techniques, but these methods are only helpful in providing broad observations. These methods could be made more useful if there were interactive interfaces that provided not only for the overviews that are currently afforded, but also for more detailed drill down to review individual submissions that appear interesting. Another technical hurdle to overcome is the need for instructors to capture thousands of submissions in some form that can be analyzed by these systems. Our experiment required significant manual effort that no instructor would be able to execute under normal circumstances.

Dynamic maps present further challenges for computationally-assisted analysis. Here we have focused on simply exploring map designs via analysis of single screen captures. The vast majority of our submissions are actually from interactive digital maps which cannot be completely summarized by a single screen capture. Therefore we see the need for new techniques to help capture and compare dynamic map projects, potentially leveraging click-stream data to assess the synthesis of map design and interaction primitives as outlined by Roth (2013).

WHAT CAN BE DONE TO AUTOMATE THE PROCESS OF DISCOVERING FRAUDULENT SUBMISSIONS?

Academic integrity issues are certainly not unique to distance learning or MOOCs, but we note here the need for better ways of discovering fraudulent submissions when faced with a massive collection of assignments to review. There is a wide range of tools available today that allow instructors to submit collections of written works to check for academic integrity violations, but to our knowledge there remains a gap in technology and services when it comes to supporting instructors who want to know whether or not a given image has been previously published. Manual searching is of course possible, but automated techniques are essential in the context of a massive course.

This problem becomes even more difficult if one considers submissions that feature interactivity, where it can be

trivially easy for students to claim another's work as their own and where similarities may not be readily detectable using image analysis methods alone.

HOW CAN QUALIFIED CARTOGRAPHIC EXPERTS BE EASILY IDENTIFIED AND ENCOURAGED TO ASSIST STUDENTS IN NEED?

Perhaps the greatest challenge we see in the further development of peer assessment techniques for cartographic education at scale is the need for map design expertise to become more scalable. While we have shown here that a relatively simple assignment with an easily understood rubric can generate consistent and reliable results, we expect grading reliability and utility to decrease as the need for detailed cartographic expertise increases for a given assignment. For example, it may be easy for students to identify the need to normalize data on a choropleth map,

but far harder to identify the incorrect use of a given projection, or the need to carefully align and distribute layout elements.

We do not, however, expect that such expertise may only reside with an instructor. Our experiences with the Maps MOOC have shown that there are significant groups of professional cartographers and geospatial analysts who take the class, even though it would seem to be far too basic for those audiences. Such students tend to be interested in trying the MOOC platform itself, and some are clearly present in order to help novices get started in cartography. It would be ideal if students with expertise were more readily identifiable such that they could be then directed by an instructor to make interventions and help solve the scalability issue when it comes to providing expert feedback.

CONCLUSIONS AND FUTURE RESEARCH

OUR WORK HERE has contributed lessons learned from the development and evaluation of peer assessment at a massive scale through experiences in teaching a MOOC on mapping. We have shown how such an assignment can be structured, what happens when students grade each other, how those grades compare to instructor grading, and how techniques from image analysis can help instructors see large collections of maps designed for a MOOC assignment. Based on these evaluations of peer assessment, we have outlined several key research challenges that require further research in order to develop mature mechanisms for evaluating map designs in massive cohorts. Our analysis of students' final map projects offers a unique evaluative approach to large map collections, assesses the extent to which students integrate theoretical concepts with current mapping tools and platforms, and can help guide future course offerings in designing content relevant to global cartographic aesthetics and demand.

As a next step in this research, we are focusing attention on the other types of feedback that we have collected from peer reviews in the Maps MOOC. In addition to numerical scores from rubric-based evaluation, most peer assessment frameworks provide for unstructured text feedback

for reviewers to explain their ratings. In the context of the Maps MOOC, these data include thousands of qualitative descriptions from peer graders, and anecdotal reports from students indicate that these explanations are critically important sources of feedback in addition to the numerical ratings. To date we have not conducted a structured analysis of these data, and we anticipate that there are more lessons to be learned from what is contained therein. Text responses on peer assessment assignments introduce another potential scale issue for educators to solve. If there are thousands of written responses, how can one instructor make sense of this feedback and use that knowledge to improve or refine a given assignment? We believe there is the potential to leverage topic-modeling tools, including methods like latent Dirichlet allocation (Blei, Ng, and Jordan 2003), to computationally extract and summarize key topics in large collections of text. These techniques are in use today for a wide range of contexts where making sense of a large corpus requires some degree of automated summarization, and initial experiments have already been conducted to explore their potential utility for analyzing the massive conversations that take place in MOOC forums (Robinson 2015).

REFERENCES

- Blei, D. M., A. Y. Ng, and M. I. Jordan. 2003. "Latent Dirichlet Allocation." *Journal of Machine Learning Research* 3: 993–1022
- Clark, A. M., J. Monk, and S. R. Yool. 2007. "GIS Pedagogy, Web-based Learning and Student Achievement." *Journal of Geography in Higher Education* 31: 225–239. doi: [10.1080/03098260601063677](https://doi.org/10.1080/03098260601063677).
- DiBiase, D., and H. J. Rademacher. 2005. "Scaling Up: Faculty Workload, Class Size, and Student Satisfaction in a Distance Learning Course on Geographic Information Science." *Journal of Geography in Higher Education* 29: 139–158. doi: [10.1080/03098260500030520](https://doi.org/10.1080/03098260500030520).
- Falchikov, N., and J. Goldfinch. 2000. "Student Peer Assessment in Higher Education: a Meta-analysis Comparing Peer and Teacher Marks." *Review of Educational Research* 70: 287–322. doi: [10.3102/00346543070003287](https://doi.org/10.3102/00346543070003287).
- Ho, A. D., J. Reich, S. Nesterko, D. T. Seaton, T. Mullaney, J. Waldo, and I. Chuang. 2014. "HarvardX and MITx: the First Year of Open Online Courses." *Harvardx and MITx Working Paper No. 1*.
- Larson, C. 2014. "Coursera's Plan for Online Education: Expansion in China." *Bloomberg Business*. Accessed June 29, 2015. <http://www.bloomberg.com/bw/articles/2014-10-27/coursera-ceo-richard-levin-plans-to-expand-the-company-in-china>.
- Luo, H., A. C. Robinson, and J.-Y. Park. 2014. "Peer Grading in a MOOC: Reliability, Validity, and Perceived Effects." *Journal of Asynchronous Learning Networks* 18: 1–14.
- McAuley, A., B. Stewart, G. Siemens, and D. Cormier. 2010. *The MOOC Model for Digital Practice*. Accessed October 27, 2015. http://www.elearnspace.org/Articles/MOOC_Final.pdf.
- McMaster, R. B., S. A. McMaster, S. Manson, and R. Skaggs. 2007. "Professional GIS Education in the United States: Models of Access and Delivery." Paper presented at the XXIII International Cartographic Conference, Moscow, Russia, August 4–10.
- Means, B., Y. Toyama, R. Murphy, M. Bakia, and K. Jones. 2010. "Evaluation of Evidence-Based Practices in Online Learning: A Meta-Analysis and Review of Online Learning Studies." Washington, DC: U.S. Department of Education, Office of Planning, Evaluation, and Policy Development.
- Nelson, J. K., and A. C. Robinson. 2015. "Understanding Map Design in the Context of a Massive Open Online Course in Cartography." Paper presented at the 27th International Cartographic Congress, Rio de Janeiro, Brazil, August 23–28.
- Nicholson, P. 2007. "A History of E-Learning." In *Computers and Education*, eds. B. Fernández-Manjón, J. Sánchez-Pérez, J. Gómez-Pulido, M. Vega-Rodríguez, and J. Bravo-Rodríguez, 1–11. Springer Netherlands. doi: [10.1007/978-1-4020-4914-9_1](https://doi.org/10.1007/978-1-4020-4914-9_1).
- Pappano, L. 2012. "The Year of the MOOC." *The New York Times*, November 2.
- Robinson, A. C. 2015. "Exploring Class Discussions from a Massive Open Online Course (MOOC) on Cartography." In *Modern Trends in Cartography*, eds. J. Brus, A. Vondrakova, and V. Vozenilek, 173–182. Springer International Publishing. doi: [10.1007/978-3-319-07926-4_14](https://doi.org/10.1007/978-3-319-07926-4_14).
- Robinson, A. C., J. Kerski, E. Long, H. Luo, D. DiBiase, and A. Lee. 2015. "Maps and the Geospatial Revolution: Teaching a Massive Open Online Course (MOOC) in Geography." *Journal of Geography in Higher Education* 39: 65–82. doi: [10.1080/03098265.2014.996850](https://doi.org/10.1080/03098265.2014.996850).
- Roth, R. E. 2013. "Interactive Maps: What we Know and What we Need to Know." *The Journal of Spatial Information Science* 6: 59–115. doi: [10.5311/JOSIS.2013.6.105](https://doi.org/10.5311/JOSIS.2013.6.105).
- Schneider, C. A., W. S. Rasband, and K. W. Eliceiri. 2012. "NIH Image to ImageJ: 25 years of Image Analysis." *Nature Methods* 9: 671–675. doi: [10.1038/nmeth.2089](https://doi.org/10.1038/nmeth.2089).

Suen, H. K. 2014. "Peer Assessment for Massive Open Online Courses (MOOCs)." *The International Review of Research in Open and Distributed Learning* 15.

Terry, J. P., and B. Poole .2012. "Providing University Education in Physical Geography Across the South Pacific Islands: Multi-modal Course Delivery and Student Grade Performance." *Journal of Geography in Higher Education* 36: 131–148. doi: [10.1080/03098265.2011.589026](https://doi.org/10.1080/03098265.2011.589026).

Unwin, D., N. Tate, K. Foote, and D. DiBiase. 2011. *Teaching Geographic Information Science and Technology in Higher Education*. New York, NY: John Wiley & Sons. doi: [10.1002/9781119950592](https://doi.org/10.1002/9781119950592).