e-ISSN 2723-9640

JoLMA
Vol. 5 — Num. 1 — April 2024

Maps and the Epistemic Risks of Visual Representation

Quill Kukla
Georgetown University, USA

Abstract Bad maps misrepresent and mislead. They hide important truths and misdirect our attention. Often, they are self-serving, promoting the values of their makers. But it is not easy to delineate what counts as a good map. Even 'good' maps that are useful, illuminating, and accurate according to their representational conventions can still mislead us, hide important patterns, and distort our understanding. In constructing a map, we necessarily balance at least three sorts of epistemic risks, which I name aesthetic risks, categorization risks, and simplification risks. Balancing these risks is always a value-laden process. Maps that employ an 'aesthetics of neutrality' can be distinctively misleading by hiding their own value-laden perspective under an aesthetic veneer of scientific objectivity.

Keywords Maps. Geographic Information Systems. Aesthetics in Science. Epistemic Risk. Representational Risk.

Summary 1 Introduction: Value-Laden Maps. – 2 Maps and Epistemic Risk. – 3 Three Types of Representational Risk. – 4 There Are No Safe Maps.

Edizioni Ca'Foscari

Peer review

Submitted 2024-01-18
Accepted 2024-02-08
Published 2024-07-26

Open access
© 2024 Kukla | CC BY 4.0

Citation Kukla, Q. (2024). "Maps and the Epistemic Risks of Visual Representation". *JoLMA*, 5(1), 39-60.

DOI 10.30687/Jolma/2723-9640/2024/01/002

1 Introduction: Value-Laden Maps

We all enjoy laughing at bad maps; there are entire Facebook groups and Twitter accounts devoted to them, with hundreds of thousands of followers. Bad maps misrepresent and mislead. They skew and hide important truths and misdirect our attention. Often, they are self-serving, promoting the values of their makers. But it is by no means easy to delineate standards for what counts as a good map, or to explain how good maps contrast with bad maps.

A seemingly simple answer is that bad maps are constructed by biased mapmakers who encode their values into their maps, whereas good maps are objective and value-neutral. But I will argue that map-making is necessarily value-driven all the way down; there is no such thing as a value-neutral map. Another seemingly simple answer is that a bad map is one that misrepresents, while a good map represents accurately. But what counts as misrepresentation? Roads are not literally black lines; the earth is not literally marked by political borders; all maps use nonliteral representational conventions. Maps are never exact copies of what they map. It is a substantial epistemological problem to demarcate the difference between nonliteral maps that serve legitimate epistemic ends and those that irresponsibly mislead. Nor does it solve this problem to say that a map must represent accurately according to established representational conventions. For one thing, as we will see, some of the best, most illuminating maps break with these representational conventions. And for another, it will turn out that some common representational conventions are especially effective at hiding the values and choices that went into the production of a map, thereby distorting our inferences from it. A *good* map, we might say, is one that generates correct and helpful inferences when it is used. But it is difficult to specify the properties that make maps good in this sense, and as we will see, even good maps also risk misleading, distorting, and obfuscating spatial knowledge.

Consider two maps that will be familiar to many in Figure 1 and Figure 2.

Figure 1 Standard Mercator map of the world

Figure 2 Washington Metropolitan Area Transit Authority subway map. 2023

Any projection map – that is, any two-dimensional map of a curved surface – cannot simultaneously preserve relative size, shape, relative distance, and direction, with respect to the shapes that it represents. Different projection systems compromise different combinations of these four parameters, but as a matter of mathematical necessity, each must compromise some of these parameters to preserve others. The Mercator map of the world **[fig. 1]**, familiar to most of us from the walls of our middle school geography classroom, preserves direction at the cost of shape and relative size. The map is accurate, according to its representational conventions. The reason this map preserves direction is because it was designed for navigational purposes. Those of us seeing the map, however, are overwhelmingly not using it to navigate, and it has well-known distorting effects upon our understanding of the world that are far from politically neutral. It centers Europe and magnifies the white-dominated global north while shrinking the global south, infamously making Greenland and Africa look comparable in size. It was also used to visually magnify the threat of communism during the Cold War (Monmonier 2018, 109). Not only do these effects scaffold racist and xenophobic narratives and understandings of space, but the original map was racist in its inception, because it was literally designed for colonizers, to help them start at the 'center' and make their way efficiently to colonizable spaces. Hence this is a *bad map* because it in effect misinforms and skews our understanding, and it does so because of poor value-laden choices in its construction. But its badness does not lie in its inaccuracy or its straying from established representational conventions.

Meanwhile, the map of the Washington, D.C. metro system **[fig. 2]**, is often heralded as an example of a good, useful map, because it serves the purposes of its intended users very well, clearly marking points of exchange, the order of stops, and city boundaries. Its bold colors and thick, elegant lines make it easy to read and use for the purpose of navigating the city. But it is in many senses a wildly inaccurate map, most noticeably in being dramatically not to scale. It is also very minimal in the information it contains. If someone were to try to use the map to find their way around D.C. by foot, or to try to extrapolate how suburbanized different areas will be, or for any number of other purposes, the map would prove misleading or useless. So, this map's value does not depend on its accuracy, or on its immunity from misleading, and this value is relative to our interest in its use.

Value-laden choices must be made throughout the course of map production. Making a map requires choosing everything from the colors and thickness of the lines, the symbology, the scale, the projection system, the data sources, the categories into which data will be divided, and the parameters the map will represent. Each choice represents and communicates some information while omitting,

distorting, or hiding other information, as cartographers themselves routinely acknowledge. Maps in their essence simplify, schematize using representational conventions, and ignore detail. This is how they communicate. A map that simply reproduced reality would not be a map at all, and given that it has to simplify and select what it shows, choices as to how to do so are inevitable. As James Scott puts the point,

> A city map that aspired to represent every traffic light, every pothole, every building, and every bush and tree in every park would threaten to become as large and as complex as the city that it depicted. And it certainly would defeat the purpose of mapping, which is to abstract and summarize. A map is an instrument designed for a purpose. We may judge that purpose noble or morally offensive, but the map itself either serves or fails to serve its intended use. (2020, 87)

In his classic book, *How to Lie with Maps*, geographer Mark Monmonier reveals the many ways in which map-making requires value-laden choices, and warns that maps that are not made with 'knowledge' and 'honesty' will be distorted and misleading. However, he does not take on the problem of what counts as a *good* map, given the ineliminability of these choice points; it is not clear what standards an honest and knowledgeable cartographer could use to achieve undistorted objectivity, or what exactly a cartographer is supposed to be honest *about*. Monmonier says that we should "be wary of cartographic manipulators" who make representational choices "that best prove their point" (2018, 159). But this is complicated by the fact that maps are communicative devices; so of course we use them to communicate our point; there is no neutral representation.

Up until World War II, geography as a discipline was focused on observing and documenting practices in local regions. It had little claim to producing generalizable knowledge and was looked down upon by the 'real' sciences. After the war, there was a revolution in academic geography, as people invented techniques for encoding elaborate statistical spatial information into maps. With the rise of GIS in the 1980s and 1990s, the capacity to do this grew enormously. Geography reinvented itself as a 'spatial science', whose central purpose was to translate quantifiable, purportedly perspective-independent and value-neutral spatial patterns and relationships in Newtonian space into objective visual representations. Some vocal geographers such as Stan Openshaw (1991) influentially argued that the only hope for geography to establish scientific bona fides was for it to become an objective and quantitative science, by focusing almost entirely on the production of such representations. The visual representation of statistical spatial information was to be the primary epistemological

method for both collecting and communicating knowledge in geography (Cresswell 2012; Pavlovskaya 2018).

The advent of 'spatial science' and GIS in particular saved dying academic geography departments, which now often earn their institutional keep by training GIS technicians, who are sought after by city planning, transportation, public health, police departments, and other such practical fields that use spatial information. The image of proper spatial science as objective and aperspectival is central to the discipline's claims to practical usefulness and scientific credibility. The institutional survival of geography departments depends on sustaining this image of objective, scientific map-making. This image sits in uneasy, unresolved, and mostly unexamined tension with the straightforward ways in which map making involves ineliminable value-laden choices.

2 Maps and Epistemic Risk

The most well-developed literature around the role of values in science concerns 'inductive risk'. The clearest definition of inductive risk for my purposes is the risk of a false negative or false positive that we accept in making a non-deductive inference from evidence to the acceptance or rejection of a general empirical conclusion from that evidence. Most paradigmatically, we assume inductive risk when we accept or reject a hypothesis on the basis of statistical evidence. This is because any time that we make an uncertain inference from empirical evidence, we assume the risk of accepting something false or rejecting something true. And, as many philosophers have shown in creative ways, we cannot decide what the correct epistemic threshold is for making such an inference – for balancing the risk of a rejecting a true hypothesis against the risk of accepting a false hypothesis – without bringing values to bear. We decide where to set this threshold in light of how bad of an outcome we think a false negative would be compared to how bad of an outcome we think a false positive could be. There is no value-independent 'right' threshold, since we are always necessarily trading one epistemic risk for another, regardless of where we set the threshold for hypothesis acceptance.[1]

But maps, like models and like other visual representations of data, are representations, not collections of inferences. They serve simultaneously as knowledge products, as evidence to be used in future inferences, and as communicative devices. The epistemic product, in

1 See the essays in Elliot, Richards 2017 for multiple explorations of this inductive risk argument; the original argument goes back to Rudner 1953 and has been much discussed in twenty-first century philosophy of science.

the case of a map, is not hypothesis rejection or acceptance, as it is in the traditional type of scientific knowledge that has been the focus in philosophy of science. It is instead a representation of information. Maps raise what Stephanie Harvard and Eric Winsberg (2021) call 'representational risks', which are epistemic risks that are distinct from inductive risks. As Harvard and Winsberg point out, the value-laden balancing of epistemic risks in representation looks different than the value-laden balancing of epistemic risks in scientific inference. When we make decisions about how to represent a state of affairs – whether through a map, a model, a taxonomy, or a visual diagram, for example – we must choose which aspects of the world to include in our representation, how to categorize those aspects, and with what symbology we will represent them. All these choices must be made in light of values and interests. Even representations that are accurate given their own conventions may mislead by encouraging incorrect inferences; by discouraging or distracting us away from important correct inferences; or by skewing our sense of salience so that we focus our inferences on the wrong things and leave important inferences unexplored. In other words, accurate representations may lead people to adopt false beliefs or to fail to adopt important true beliefs.

As Harvard and Winsberg point out, representations are not themselves true or false. Representations do not make claims, but instead they lead us well or poorly in our inferences and belief formation. The risk is not that a representation will turn out to be false, but rather than it may distort the reasoning of those who use it. Because *all* representations are partial and involve choices about how and what to represent, all representations come with this representational risk of misleading. Moreover, because representations are inherently *communicative*, the risk that they mislead is a risk that should be considered internal to the process of representing, and not just a separate piece of moral luck. Methodological choices in representation are inextricable from communicative choices.

3 Three Types of Representational Risk

Maps raise representational risks of at least three kinds, corresponding to three necessary stages in the map-making process. I name these aesthetic risk, categorization risk, and simplification risk. Maps are inherently aesthetic, categorizing, and simplifying. In no case can any of these three types of risk be avoided; they can only be managed in light of our values and interests, which govern how we want the map to be used and what sort of knowledge we want it to convey. In each of the three cases, there is no 'safe' or neutral answer to the question of how to manage the epistemic risks involved.

a. *Aesthetic risk*: Maps are perceivable objects, and they are inseparable from their own aesthetic features. Aesthetic risk is the epistemic risk introduced by making necessary *aesthetic choices* about representational conventions, which will necessarily make some information salient at the cost of other information, and shape users' attention to and interpretation of the map. There is no 'true' or 'neutral' set of aesthetic conventions; all aesthetic choices must be made on the basis of values and purposes.
b. *Categorization risk*: Of necessity, maps introduce representational conventions where symbols stand for categories of objects or data. For instance, solid black lines may stand for all roads with open intersections, while double lines may stand for all controlled access freeways; this convention groups roads together by how they are accessed. A map that represents median income by neighbourhood may divide income by quartiles, quintiles, standard deviations, or infinite other ways, but since it cannot represent each of the infinitely many possible incomes a different way, it must pick some categorization system or other. Reality does not hand itself to us pre-categorized, and so these categorization choices must be made in light of values and interests. In turn, these choices shape what patterns the map reveals and what patterns it hides, which in turn directs how people make inferences from the map.
c. *Simplification risk:* Maps do not represent spatial reality in all its complexity, but rather selected parts of this complexity. Mapmakers must choose which features, parameters, and relationships to include on their map. Maps that include more details and parameters have more nuance and contain more information, but they risk communicating less as they become visually incomprehensible. Communicative trade-offs of this sort are inevitable; leaving off a parameter or feature risks making it invisible and unsalient to readers of the map, while including extra features clutters the map and reduces its communicative power. Either way, the map can mislead.
d. Aesthetic choices and choices about categorization and simplification are essential, unavoidable parts of map-making. In none of the three cases is there a value-neutral answer written into nature as to how these choices should be made. In the following three sections, I explore how complex the role of values in making all three sorts of choices can be. Maps that are accurate, follow established representational conventions, and communicate important truths still have the real potential to mislead because of these choices.

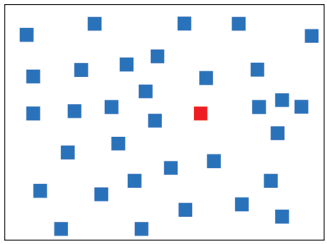

Figure 3 Find the red square

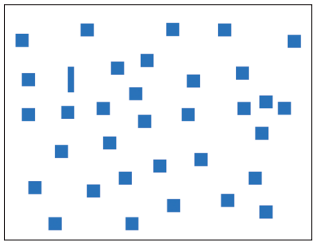

Figure 4 Find the blue line

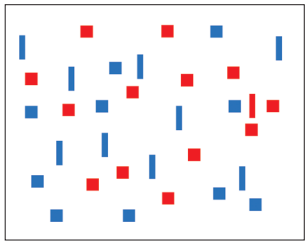

Figure 5 Find the red line

A. Aesthetic Risk

Producing a visual representation of spatial information is never a matter of mechanically processing data. Rather, it requires making aesthetic choices. Unlike when we represent knowledge or evidence in propositions, maps necessarily have a visual, aesthetic form. In producing a map, aesthetic choices are not external to the content, like picking the font for a book; they are part of what individuates the map. There is no abstract map independent of its symbology: the colors it uses, the thickness of its lines, the contrast between figure and ground, the shape of its nodes. These choices directly affect what information the map makes salient and how the viewer will interpret and use it. For example, contrasting colors, up to a point, make information stand out and allow for quicker cognitive processing. Thus, color contrast can be used to shape what the map makes salient and what it hides (Fu et al. 2013). But too many colors become difficult to process, and have the opposite effect, lessening how much information the map effectively communicates. A map literally *tells* us different things depending on these kinds of aesthetic choices, even holding the information strictly contained in the map constant.

It is easy to confirm phenomenologically that aesthetic choices directly impact how we process and extract information from visual representations. Below is a small experiment that you can perform for yourself while reading this paper. First, look at **[fig. 3]** and find the red square as fast as you can. Next, look at **[fig. 4]**, and find the blue line. Finally, look at **[fig. 5]**, and find the red line. All three tasks are easy, but I trust it is subjectively obvious how the processing time increases across the three tasks.

This experiment directly demonstrates how aesthetic features like color and shape can impact our perceptions of salience and how we process the information in a visual representation. For this reason, two maps can encode exactly the same spatial information and be technically exactly equally accurate in doing so, and yet they may communicate very differently because of different aesthetic choices

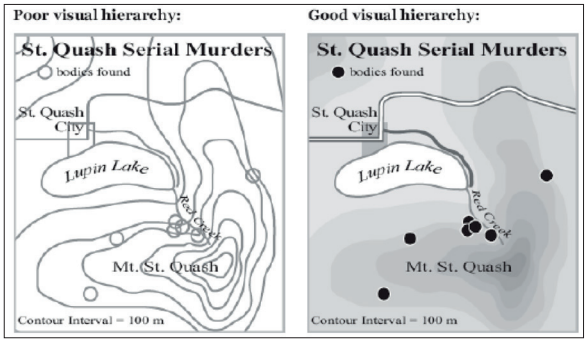

Figure 6
Two maps with identical spatial information and accuracy

in their construction. See for example **[fig. 6]**, which shows two informationally identical maps of murder sites.

There is no objectively correct aesthetic look that a map should have; which aesthetic choices we should make depends upon which patterns we are trying to make salient. The map on the right in **[fig. 6]** is 'better' because it better serves the purpose of making the murder sites and their topographic context salient and easy to process, and this is likely what the map is used for. But there are infinite other patterns consistent with the information encoded that it does not make salient. Any set of aesthetic choices comes with *aesthetic risk* – the risk that aesthetic choices that make some patterns salient and enable some inferences will hide other patterns, in ways that can mislead. And this risk cannot be managed in a value-free way.

An exhibit at the Pratt Institute in October 2017, entitled *You Are Here NYC: Art, Information, and Mapping*, featured artworks that used geospatial data to produce representations designed to give aesthetic insight into New York City as a human place. It included Doug McCune's piece, *Data Sketch: Routes*, which used GIS data to create a map of New York City with a three-dimensional double wall around its boundaries. The height of the taller wall indicates the number of immigrants to the city who arrived from that compass direction, while the height of the second, surrounding wall indicates how many of these immigrants were children **[fig. 7]**.

It also included Xingying Du, Michelle Htar, and Jessica Silverman's *Journeys Disconnected – Reconnected*, which used colored yarn to track the migration patterns of 66 people who were buried at Hart Island in the Bronx, which served as a prison, a psychiatric institution, a sanatorium, and a boys' reformatory. The piece makes vivid how the island served as a terminus for complex global stories **[fig. 8]**.

Such artworks are quantitatively representations of spatial data, according to their representational conventions. The main thing that

Figure 7 Doug McCune, *Data Sketch: Routes*. 2017

Figure 8 Xingying Du, Michelle Htar, Jessica Silverman, *Journeys Disconnected – Reconnected*. 2017

distinguishes them from traditional GIS maps is that their makers are self-reflective about their aesthetic choices in producing such representations, and about how the representations are designed to make certain spatial patterns with human significance salient. Their aesthetic choices are specifically designed to encourage specific interpretations of and responses to the spatial information they portray. The aesthetic features of these works are not overlain on top of regular maps, nor are these artworks instead of regular maps. Instead, both reveal the extent to which every map is a product of necessary value-driven aesthetic choice points.

A more standard immigration map has the look of objectivity rather than art. However, its makers also had to choose its color scheme and symbology. The choice to make a map with the look of stark neutrality – one that seems to depend on no investments in a specific interpretation, and no interest in capturing any particular lived perspective – is itself a value-laden aesthetic choice. The aesthetics of withholding any interpretive perspective is a look that we choose because of specific interests. It is no more or less epistemically fecund or representationally accurate in virtue of this aesthetics of neutrality. In no obvious sense can it be said to have any special 'epistemological priority' over maps that make more explicitly communicative aesthetic choices. Whether or not the makers of a 'neutral looking' immigration map **[fig. 9]** were deliberate and thoughtful about the downstream effects of their aesthetic choices, they had to choose all the same, and their choices affect what the map communicates, and which inferences it stimulates and which it discourages.

All three maps are good maps for some purposes. All of them are, as far as we can tell, accurate given their own representational conventions and goals, and they make salient important truths about migration, borders, and the kind of place that New York City is. McCune's map is based on quantitative data, but it conveys a qualitative

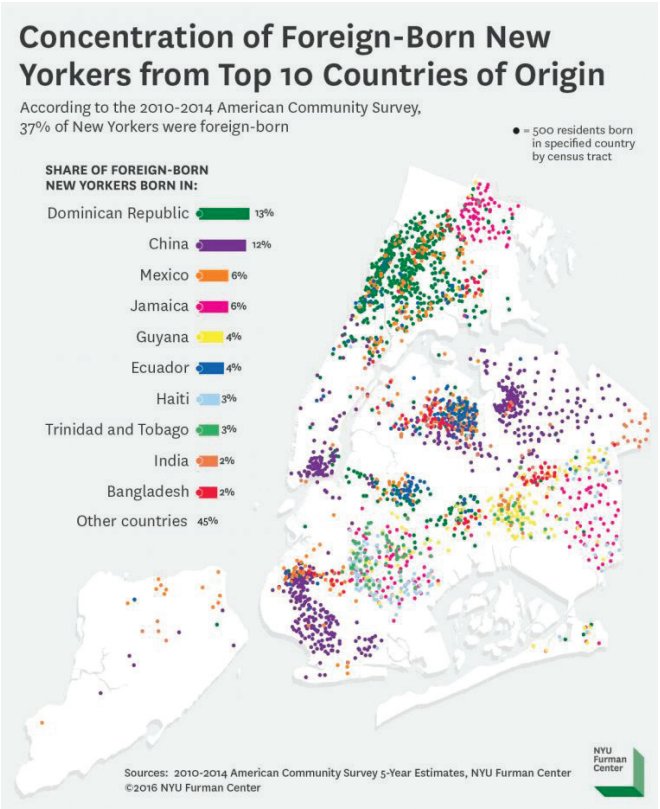

Figure 9
Furman Center immigration map of New York City. 2015

sense of place and migration that are absent in the more 'objective'-looking map. Du, Htar, and Silverman's map brings spatial paths traversed to life as lived journeys. The Furman Center map is good at conveying where there are ethnic neighborhoods in the city, while it evokes no sense of place or journey. We might think that because it is an 'objective' map, it is not its job to convey anything so humanistic as a sense of place or journey. But the other two maps prove that representationally accurate maps can do this. Whether they do or not depends not on the objective correctness of the map, but on the aesthetic choices made during its construction.

All three maps make different value-laden choices based on which patterns they care about conveying and what sorts of inferences they care about stimulating. All three maps have legitimate epistemic value. All three maps balance aesthetic risks, since their aesthetic choices make some meanings and patterns salient while thereby directing attention away from others.

B. Categorization Risk

A map will not be readable unless it categorizes its data. This might be as simple as two-lane roads being symbolized in one way and highways in another. Or it might be a matter of breaking data for income, temperature, or some other continuous variable into quartiles, quintiles, or deciles on a choropleth map. Or it might involve breaking up data using racial or gender categories. These categorizations are necessary, because not every data point can have a unique symbol, otherwise the map will not reveal spatial patterns any more vividly than a photograph. But these categorizations are choices, which are not written into nature. Categorization choices are not true or false or even accurate and inaccurate, but rather adequate or inadequate to purpose. How a map maker chooses to categorize their data will directly affect which patterns the map reveals and which it hides.

Mark Monmonier writes,

> A single set of numerical data can yield markedly dissimilar maps. By manipulating breaks between categories of data to be shaded on a choropleth map, for instance, a mapmaker can often create two distinctly different spatial patterns... Wary map users must watch out for statistical maps carefully contrived to prove the points of self-promoting scientists, manipulating politicians, misleading advertisers, and other propagandists. (2018, 153)

He warns about unscrupulous and irresponsible mapmakers bending maps to their own ends through categorization systems, but there is no correct categorization system built into nature. All of them come with representational risk, because all hide some patterns while revealing others. His warning about manipulative categorization systems suggests that there exist, in contrast, honest and disinterested categorization systems, but it is unclear what these would be or by what standards we would determine them.

For a simple example, consider the two maps of American mpox infections from 2022 **[figs 10a-b]**. These maps were based on the same data set and represent the same time. Such maps typically use more intense colors to indicate locations with higher infection rates. But of course, we can make mpox look like an especially pressing issue or a relatively mild concern depending where we set the cut-off for a 'very high' infection rate, which varies dramatically on different maps. Notice how much more pressing the issue looks on one of these maps than the other. Notice also that we can see different patterns in each map.

Consider now a map showing the residence locations of people of different races. Such a map requires that we decide how to categorize people by race. This means choosing a racial categorization system.

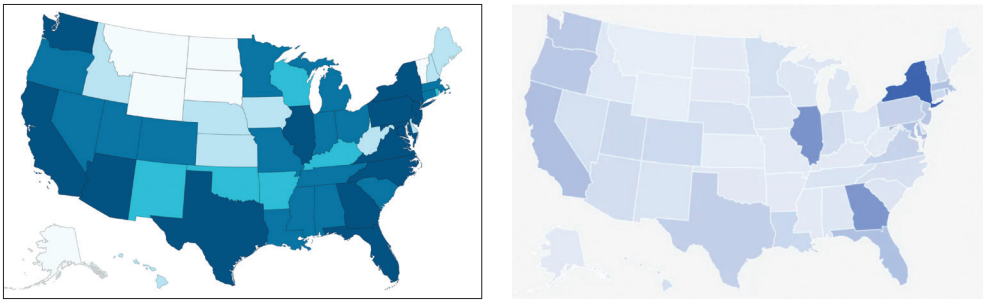

Figure 10a-b Mpox infections by American state. 2022

[Fig. 11] shows a dot map of Detroit, produced by Dustin Cabel at the University of Virginia. Cabel used 2010 census data to represent the racial demographics of each major American city, with one racially color-coded dot for every person deemed a resident.

In this map, one can vividly see the stark segregation of the city, with 8 Mile Road a sharp dividing line between White and Black. This map is an excellent map, in that it uses categorization and symbology to reveal an important truth about the city of Detroit accurately and vividly; it makes a pattern lucid for us in a way that a list of statistics never could. When you look at this map, the boundary created by 8 Mile Road takes on hard reality. We can *see* segregation at work! This map encourages correct and important inferences that would not otherwise be salient.

But the map is a product of a series of categorization choices that hide other patterns. For instance, following the census, the project folded all Middle Eastern and Arab people into the 'White' category. It thereby made it impossible to see one of the most vulnerable and vexed groups in the country. Indeed, the Detroit greater area has among the largest Muslim populations in the United States (or the largest, depending upon the source and the exact definition of the metro area). This is an important fact about the spatialized racial politics of the city, which is obscured by the map. The map also folds together all Asians, despite important differences between groups from different parts of the continent. For instance, this means that this map does not reveal the distinctive physical isolation of Detroit's Hmong community, which again is an important feature of the spatialized racial politics of the city (Yang 2003).

So, it is an objective flaw in the map that it used categorization systems that occluded these spatial patterns? No, because using a more fine-grained racial categorization system would have muddied the pattern that did emerge, which is a real and important pattern. The map as it stands is not wrong or flawed, but what it reveals and

Figure 11
Central Washington, D.C. commuting map, 2019

what it hides is based on value-laden choices that vanish under the veneer of stark objectivity that the visual representation suggests.

The racial categorizations used to produce this map were inherited pre-packaged from those used in the United States census. We do not know whether the cartographers thought explicitly about the impact that using the census categorizations would have on these dot maps. Making use of a secondary source of data like the census provides concrete benefits in map-making, giving cartographers access to vast data sets that would be impossible for individual researchers or small research teams to recreate. Hence some value-laden choices were already baked into the technological infrastructure available to the researchers who produced the maps. Thus, the researchers made a second kind of epistemic trade-off: they accessed a larger data set than any they could produce on their own, at the cost of offloading the burden of reflecting on racial categorizations and the epistemic risks they pose.

The visual representation that results from these choices affects what we see as real boundaries and divisions in the city. In turn, this affects not just what theoretical inferences we draw, but also our practical decisions. Maps like this one influence investors' and developers' choices about where to buy and build property, individuals' choices about where to live and visit; and policy decisions concerning transportation infrastructure and the like. Thus, categorization choices that determine which patterns our maps reveal and which they hide concretely impact the world we live in. The map in **[fig. 11]** is, by any reasonable measure, a good map that reveals useful and important truths, but it is not free of representational risk. We need to read it with a critical eye, aware of what sorts of patterns it may be occluding.

C. Simplification Risk

Maps necessarily simplify the spatial complexity of the area that they map; any map must select which spatial features, parameters, and relationships it will represent. In the last section, we saw that we need to choose how to sort individual data points into categories, which is also a form of simplification. But before we even get to categorization, we need to pick what will be represented in a map at all, abstracting away from the indefinitely rich particularity of any space and choosing parameters to represent. Including more parameters on a map increases nuance and detail, but lowers its communicative power and inferential ease.

There has been a strong and unsurprising bias, within GIS culture, in favor of representing simple quantitative relationships and features. GIS maps achieve their veneer of scientific legitimacy and objectivity partly by sticking to measurable, visually clean spatial relationships and data points. There is also a bias in favor of including parameters for which we have large data sets, as this increases the statistical power of the map. But here, representational risk comes in, because both biases mean that we are less likely to represent and communicate more complex relationships and non-quantitative spatial phenomena, which means that those patterns are occluded, and don't show up as part of our visual canon of objective spatial truths. In other words, what it is easiest to map effectively unsurprisingly impacts what we map, and in turn this shapes our understanding of spatial reality, with the easily mappable parts showing up as objective features of the world, while other patterns are left invisible.

Consider, for example, commuting maps. Geographers love to make maps that represent commuting patterns (that is, trips between home and work), perhaps divided by gender, or race, or by income bracket. It is worth exploring why there are so many commuting maps, and which patterns and relationships are revealed and hidden when we use GIS to represent commuting patterns. I want to highlight what commuting maps tell us about *what kinds of motion* and *whose motion* are essential to understanding the dynamics of a region.

Why do our maps of motion through cities so often focus on home-work trips? First, we tend to think of these as the 'main' trips that define someone's day. Second, we have the best data sets for this kind of motion. Commuting data is easily available, because we keep track of where people live and where they work, but we do not have any immediate access to the other ways in which people move through space. These are connected facts: this data is more easily available partially because we take home and work to be the two main places where people belong, the places it is worth collecting data about. In contrast, we cannot easily get data to map movement through space to visit family, care for a parent, go to church, pick children up through

school, go to the doctor, socialize, and so forth. In fact, only about 25% of vehicle trips in the United States are commuting trips, according to a 2019 study (Tsafos 2019). That figure predates the COVID pandemic, so it is likely that the current number is even smaller in many cities, because of the large rise in the number of people who work remotely from home. Thus, a commuting map will give a highly distorted picture if we see it as a representation of city movement.

Whose lives are centered and privileged by these maps, and what aspects of those lives are taken as essential? First, it centers people who have one single work location, as opposed to those who have other kinds of jobs or who are not traditionally employed. This leaves out artists, contract workers, adjuncts teaching at multiple universities, many sorts of tradespeople such as plumbers and painters. It privileges white-collar in-person office workers and blue-collar manufacturing workers. It leaves out very young and very old people and many disabled people. It is more likely to leave out women than men. Meanwhile, policy and transportation decisions, such as where to place bus routes and bike lanes, are made on the basis of these maps, which in turn makes commuting trips easier, but often makes other sorts of motion through a city more difficult. Thus, such maps encode whose lives and what dimensions of those lives are privileged, while in turn contributing to the development of infrastructure that further privileges those lives and life dimensions.

Consider the map of (pre-COVID) commuting patterns in central Washington, D.C. **[fig. 12]**.

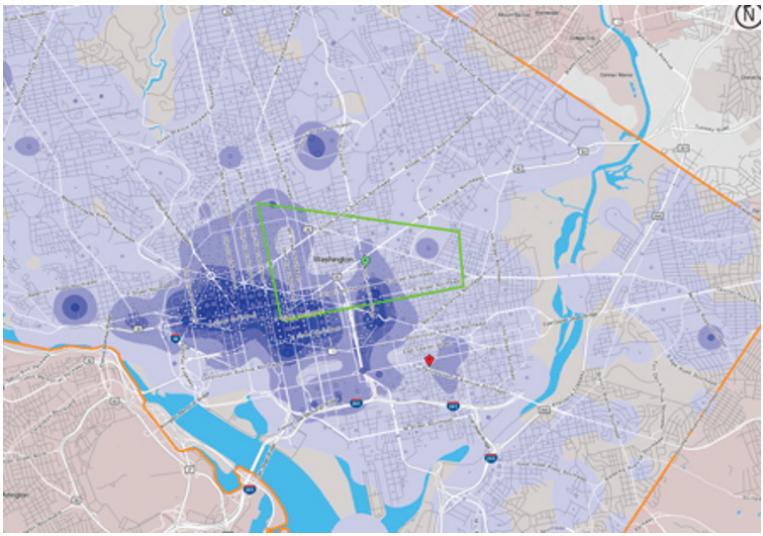

Figure 12 Central Washington, DC commuting map. 2019

This is a map that I generated using census data, which shows the workplaces of people who live in the fastest growing, most quickly gentrifying part of Washington, D.C. (inside the central polygon). Looking at the map, you can see a clear commuting direction: people who live in the center of the city tend to travel to the west and southwest on their commute. Most people living in this central polygon are young white educated professionals who have recently migrated to the city. For anyone familiar with Washington, D.C., this map will not be surprising; the dark areas representing heavy employment for those living within this intensely gentrified area include medical facilities and governmental organizations such as the World Bank in Foggy Bottom, and political offices and lobbying and law firms along K Street NW, as well as Georgetown University and George Washington University.

I originally created this map for the purpose of showing why people were advocating for building bike lanes leading from the central polygon to the areas to the west and southwest, and the map is effective at demonstrating this need. Since many young professionals move to the city partly to rid themselves of car dependence, and since biking is a healthy, cheap, and environmentally sustainable mode of transportation, the case for bike lanes appears transparent from this map. But there is a second story that this map not only fails to tell, but occludes. The area in the center of the city that shows up as to be commuted through on this map is also one which elderly Black residents, who have managed to stay in their homes as the area gentrifies, move within in order to go to one of several local, longstanding Black churches and to visit family and neighbors. It is also an area that Black people displaced by gentrification, who have moved out to public transportation-starved, car-dependent Prince George's County, Maryland, travel into, to see elderly relatives and to go to church. None of this motion shows up on the map.

Thus, this map makes the area look like one to be passed through by residents from the east on their way to work in the west/southwest, rather than one to come into or move around in. The map gives an apparent place-meaning to the space that we are seeing, and in turn this can shape decisions about what sorts of infrastructure the area needs. It turns out that bike lanes through this area will disrupt much-needed parking for Black churches and other key Black community institutions. This commuting map, if taken on its own, thus encourages further gentrification and disruption of the at-risk ecology of the area.

What I wanted to do, in principle, was to make maps showing both kinds of motion, so that I could demonstrate the nuanced and contradictory mobility needs of the neighborhood. But this proved impractical, because the GIS technology and data available to me through the census privileged home-work trips. Showing the uses of the space by Black residents would have required that I somehow collect information on movement around and into the neighborhood

person by person, and the results would have been unreliable and have had much less statistical power than the giant census data sets. Our privileging of some kinds of motion over others is built into our institutions and technology, and reflected in our maps. This commuting map is not inaccurate; it represents what it claims to represent. But it is difficult, when looking at it, not to see commuting as a proxy for motion in general through the city. This is a representational risk that the map poses. Based on such commuting maps, we understand the city as a set of flows that require various kinds of infrastructural support, thereby hiding the lives, motions, values, and needs of those whose spatial lives are not organized around a traditional home-work commute. A mismatch between what this map encodes and what it communicates is likely, given the background assumptions that frame its look and its uptake. It lends itself to misuse, if this misuse is not actively blocked by further communication and framing, or by a countermap that is hard to produce. Here again we have a map that is accurate according to relatively straightforward representational conventions and can communicate important and useful truths about spatial patterns. But it still runs the risk of misleading its users and occluding important information.

4 There Are No Safe Maps

I have tried to show that any map production is riddled with ineliminable representational risks, including what I have called aesthetic, categorization, and simplification risks, so that the values we use to manage these risks will always affect which real patterns the map communicates and which it hides. There are no neutral, value-free choices to be made about aesthetics, categorization, or simplification. Once we recognize that every map occludes and reveals, and that value-laden choices at multiple stages determine how it does this, we can lose our handle on what counts as a good map. A good map, one would think, should be epistemically helpful and fecund, rather than epistemically inhibiting. But I have argued that even well-constructed, epistemically helpful and fecund maps can also be epistemically inhibiting. There does not seem to be any straightforward way to classify some maps as epistemically safe. A map that is designed to direct people's attention in ways that are misleading and epistemically damaging is a bad map, and this is true even if it is accurate according to its representational conventions. But as we have seen, many maps that are well-designed and epistemically illuminating will at the very same time leave out and occlude other important spatial patterns, and can thereby do epistemic damage.

One reason why many maps are likely to occlude important patterns for their users is that they often follow conventions that give

them the look of neutrality and transparency. Many maps present visual data in a way that makes us feel like we are seeing straight through to the objective truth about spatial data. This is partly simply because we find visual representations powerful, but it is also because we have entrenched a specific conventional aesthetic and conventional choices about what and how we map that signal scientific accuracy and value-neutrality in our maps (Ferdinand 2019). Some maps look business-like and neutral, like the map in **[fig. 9]**. Maps that use what we read as a 'neutral' aesthetic, and that represent simple quantitative relationships, look the most objective and transparent to us. Maps that present the world as *neutrally* calculable, measurable, statistically analyzable, and representable are prized, for reasons that we saw at the start of this presentation; they give geography its institutional claim to legitimacy and scientific bona fides. However, they also hide their own representational risks and discourage critical distance from their message. Their apparent transparent objectivity is epistemically misleading. Arguably, maps that aesthetically display their own specific perspective and value-laden production and perspective have distinctive epistemic advantages. For instance, the 'artistic' maps we saw in **[figs 7-8]** had an evident point of view that they were trying to communicate. They seem more honest in their open use of aesthetic choices than do many more 'neutral' looking maps, and they are also informative, even though they too come with representational risks.

Critical GIS is a small but flourishing area devoted to revealing how GIS conventions in geography are not value-neutral. Critical GIS theorists and practitioners push back against the institutional privileging of statistics-rich maps that represent the kinds of relationships that are easy to quantify, and that present themselves as value-free. These maps systematically occlude other sorts of spatial patterns, including patterns that could only be revealed through qualitative inquiry. Critical GIS theorists and practitioners promote uses of GIS that are openly governed by value-laden goals such as social transformation, and the discovery of marginalized and hidden spatial patterns whose visibility has political value (Pavlovskaya 2006; 2018). Often, this involves finding ways to represent qualitative data in map form.

For instance, in 2008, Mei-Po Kwan conducted an already-classic study of Muslim-American women's felt danger and safety as they moved around Detroit in the wake of September 11. She used GIS technology to create phenomenological maps. Each map required the curation of individual qualitative data. Such qualitative maps take a lot of labor to produce. Her maps and their framing reveal clearly that Kwan is invested in unearthing a specific spatial perspective that she takes to have value; she is not revealing the neutral truth about the organization of space. **[Fig. 13]** represents one Muslim woman's phenomenological map of her experiences of safety and danger

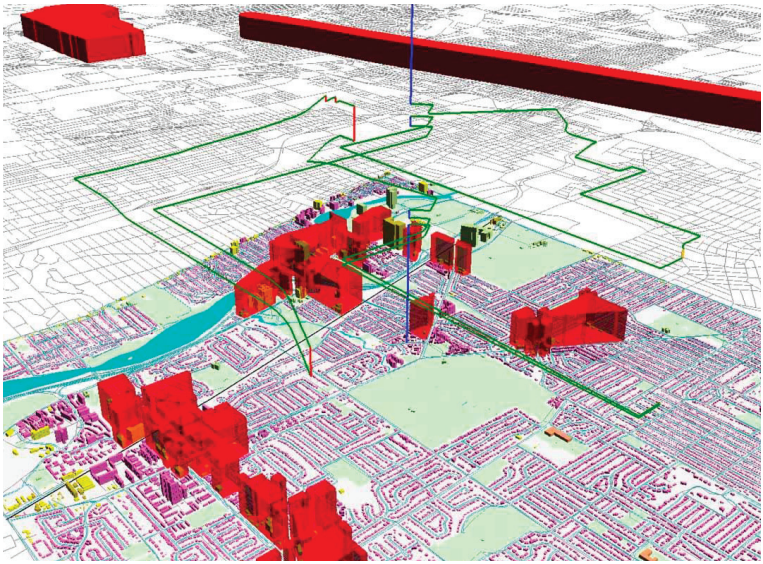

Figure 13 From Mei-Po Kwan 2008

as she navigated her city after September 11. Green lines represent her motion through the city, and red patches represent spaces she experienced as unsafe.

Such maps, like any maps, balance epistemic risks. They are the product of aesthetic, categorization, and simplification choices. They sacrifice some epistemic values, such as statistical power, in order to promote others, such as revealing experienced spatial patterns and phenomena that are hidden by standard quantitative methods. Of course, these sorts of openly value-laden and qualitative maps cannot supplant traditional, conventional maps, as there are many relevant and important patterns that they do not convey. But they have special epistemic value insofar as they reveal real patterns without encouraging the viewer to see these patterns as neutral. They are not pretending to value-neutrality or aperspectivality, but rather trying to capture and communicate a specific perspective accurately in visual, spatial form.

Mapmaking requires a dense and ineliminable series of value-laden choices that balance epistemic risks; there are no such things as safe or neutral maps. The resulting representations are communicative tools that always have the potential to mislead and misdirect attention, as they always hide some patterns in order to reveal others. Maps that use tropes and representational conventions that suggest their own neutrality and transparency may be distinctively

misleading, since they do not invite reflection upon the values and purposes that guided their creation, or upon on the patterns that they occlude. There is some irony in this, since these are the maps most highly prized as science rather than art.

References

Cresswell, T. (2012). *Geographic Thought: A Critical Introduction*. Hoboken, NJ: John Wiley & Sons.

Elliott, K.C; Richards, T. (eds) (2017). *Exploring Inductive Risk: Case Studies of Values in Science*. Oxford: Oxford University Press.

Ferdinand, S. (2019). *Mapping Beyond Measure: Art, Cartography, and the Space of Global Modernity*. Lincoln, NE: University of Nebraska Press.

Fu, K. et al. (2013). "Superpixel Based Color Contrast and Color Distribution Driven Salient Object Detection". *Signal Processing: Image Communication*, 28(19), 1448-63.

Harvard, S.; Winsberg, E. (2022). "The Epistemic Risk in Representation". *Kennedy Institute of Ethics Journal*, 32(1), 1-31.

Kwan, M.-P. (2008). "From Oral Histories to Visual Narratives: Re-Presenting the Post-September 11 Experiences of the Muslim Women in the USA". *Social & Cultural Geography*, 9(6), 653-69.

Monmonier, M. (2018). *How to Lie With Maps*. Chicago: University of Chicago Press.

Openshaw, S. (1991). "A View on the GIS Crisis in Geography, Or, Using GIS to Put Humpty-Dumpty Back Together Again". *Environment and Planning A*, 23(5), 621-8.

Pavlovskaya, M. (2006). "Theorizing With GIS: A Tool for Critical Geographies?". *Environment and Planning A*, 38(11), 2003-20.

Pavlovskaya, M. (2018). "Critical GIS As A Tool For Social Transformation". *The Canadian Geographer/Le Géographe Canadien*, 62(1), 40-54.

Rudner, R. (1953). "The Scientist Qua Scientist Makes Value Judgments". *Philosophy of Science*, 20(1), 1-6.

Scott, J.C. (2020). *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. Yale: Yale University Press.

Tsafos, N. (2019). "The Slowly Changing U.S. Commute". *Center for Strategic and International Studies*, June 12. https://www.csis.org/blogs/energy-headlines-versus-trendlines/slowly-changing-us-commute.

Yang, K. (2003). "Hmong Diaspora of the Post-War Period". *Asian and Pacific Migration Journal*, 12(3), 271-300.