Tutorial

Effective Maps, Easily Done: Visualizing Geo-Psychological Differences Using Distance Weights

Tobias Ebert[1] ![ORCID], **Lars Mewes[2], Friedrich M. Götz[3], and Thomas Brenner[4]**

[1]Mannheim Centre for European Social Research, University of Mannheim, Mannheim, Germany;
[2]Institute of Economic and Cultural Geography, Leibniz University Hannover, Hannover, Germany;
[3]Department of Psychology, University of British Columbia, Vancouver, BC, Canada; and [4]Department
of Geography, Philipps-University Marburg, Marburg, Germany

Abstract

Psychologists of many subfields are becoming increasingly interested in the geographical distribution of psychological phenomena. An integral part of this new stream of geo-psychological studies is to visualize spatial distributions of psychological phenomena in maps. However, most psychologists are not trained in visualizing spatial data. As a result, almost all existing geo-psychological studies rely on the most basic mapping technique: color-coding disaggregated data (i.e., grouping individuals into predefined spatial units and then mapping out average scores across these spatial units). Although this basic mapping technique is not wrong, it often leaves unleveraged potential to effectively visualize spatial patterns. The aim of this tutorial is to introduce psychologists to an alternative, easy-to-use mapping technique: distance-based weighting (i.e., calculating area estimates that represent distance-weighted averages of all measurement locations). We outline the basic idea of distance-based weighting and explain how to implement this technique so that it is effective for geo-psychological research. Using large-scale mental-health data from the United States ($N = 2,058,249$), we empirically demonstrate how distance-based weighting may complement the commonly used basic mapping technique. We provide fully annotated R code and open access to all data used in our analyses.

Psychologists have long been interested in the psychological and cultural characteristics of places (Oishi & Graham, 2010; Rentfrow & Jokela, 2016). However, it is only now—spurred by the rapidly expanding availability of large-scale data—that research examining geographical differences in psychological phenomena is really booming (Ebert et al., 2022; Hehman et al., 2020; Rentfrow, 2020). The topics and data studied under this new umbrella are diverse and span multiple psychological subdisciplines. To illustrate, researchers have used large-scale online personality surveys (Chopik & Motyl, 2017; Elleman et al., 2020; Rentfrow et al., 2008, 2013), large-scale experimental bias assessments (Hehman et al., 2019; Ofosu et al., 2019), or massive amounts of geotagged digital footprints (Eichstaedt et al., 2015; Jaidka et al., 2020; Obschonka et al., 2019) to examine spatial variation in psychological phenomena.

One, perhaps the most, compelling way to describe and communicate the spatial organization of psychological phenomena is to visualize it in maps. Indeed, maps are a central element of almost every article that investigates

Corresponding Author:
Tobias Ebert, Mannheim Centre for European Social Research,
University of Mannheim, Mannheim, Germany
Email: tobias.ebert@uni-mannheim.de

geo-psychological differences (Ebert, Götz et al., 2022). However, mapping techniques are typically not a standard feature of psychologists' methodological training. As a result, most published geo-psychological studies rely on the most basic mapping technique: color-coding disaggregated data (i.e., grouping individuals into pre-defined spatial units and then mapping out average scores across these spatial units). To be clear, there is nothing wrong with this mapping technique. However, as we outline in this article, in many research contexts, this basic mapping technique may leave some unleveraged potential to effectively visualize spatial data. Therefore, here we introduce and implement an alternative technique to visualize spatial data: distance-based weighting (i.e., calculating area estimates that represent a distance-weighted average of all measurement locations). We believe that distance-based weighting requires only little incremental effort but may offer large incremental utility in many research contexts. Our aim is to enable readers to readily implement distance-based weighting in their workflow. We provide fully annotated code to reproduce all analytical steps in R.

Demonstration Data: Mental Health in the United States

We demonstrate all empirical steps using real-world data. To do so, we chose (a) a prominent psychological attribute (i.e., mental health) that most readers will likely be familiar with and (b) a data set that mirrors common geo-psychological research in size and spatial-coverage characteristics. The following files are necessary to replicate our analysis and can be downloaded from our OSF project page (https://osf.io/quhcb):

> mental_health_counties.csv: contains information on the average mental health (based on a sample of 2,058,249 U.S. residents [*observations*]) of a county's population[1] (*menthlth*: Now thinking about your mental health, which includes stress, depression, and problems with emotions, for how many days during the past 30 days was your mental health not good? [count variable range = 0–30]). These data pool the survey years 2005 to 2010 of the Behavioral Risk Factor Surveillance System (Centers for Disease Control and Prevention [CDC], 2020).[2]

> county_shape.shp and state_shape.shp with corresponding .prj, .dbf, and .shx files: These so-called shapefiles originate from the U.S. Census Bureau (2020). They contain digital vectors to store polygons representing the spatial location and boundaries of the 3,108 counties and the 48 states of the contiguous United States plus Washington, D.C.

The following code imports, prepares, and merges these data files:

```
# Import spatial data
county_layer <- read_sf("us_counties
  .shp")
state_layer <- read_sf("us_states
  .shp")

# Import county-level variable
my_data <- read_csv("mental_health_
  counties.csv")

# Select relevant columns
sp_var <- "menthlth"
nobs_var <- "observations"
geo_var <- "fips"
vars <- c(geo_var, sp_var, nobs_var)
my_data <- my_data %>%
  select(all_of(vars))

# Rename columns to make code generic
colnames(my_data) <- c("geo_id", "sp_
  var", "nobs")

# Join spatial data with county-level
  data
my_data$geo_id <- as.character(my_
  data$geo_id)
county <- left_join(county_layer, my_
  data, by = c("GEOID" = "geo_id"))
```

Conventional Mapping

First, we briefly present the basic mapping technique that underlies almost all existing geo-psychological studies, that is, color-coding disaggregated data. This means that for each geographical unit, the corresponding average score of the phenomena under study is mapped out. In our case, we would (a) calculate the average number of poor-mental-health days per month for each county and then (b) map out this average county by county. Note that we can map out average scores only for counties whose samples are sufficiently large to derive meaningful estimates. Previous research has used varying minimum inclusion thresholds, such as 50 observations (Götz et al., 2018; Matz & Gladstone, 2020), 100 observations (Ebert et al., 2019; Obschonka et al., 2018; Payne et al., 2019), 150 observations (Hehman et al., 2017), and 200 observations (Bleidorn et al., 2016) per geographical unit. We here sought to choose a cutoff that is sufficient to produce reliable aggregate-level health scores and allows us to retain a large number of
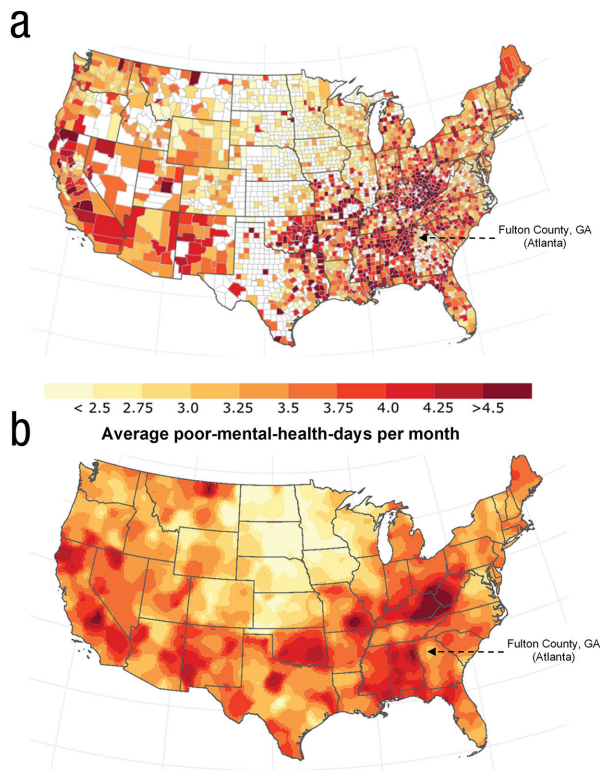
a

< 2.5 2.75 3.0 3.25 3.5 3.75 4.0 4.25 >4.5

Average poor-mental-health-days per month

b

Fig. 1. (a) Conventional mapping of regional differences in mental health across 3,108 U.S. counties. White colored counties did not meet the minimum inclusion threshold. (b) Map of regional differences in mental health across the United States using spatial grids and distance weights.

```
# Reduce sample to counties with at
  least 50 participants
disaggr_layer <- county %>%
  filter(nobs >= 50)

# Define color coding scheme and
  assign colors
grouping <- seq(2.5, 4.5, by = .25)
colors <- brewer.ylorrd(10)
disaggr_layer$group <- findInterval
  (disaggr_layer$sp_var,
  grouping) + 1
disaggr_layer$colors <- colors
  [disaggr_layer$group]

# Map mental health across counties
  while depicting state boundaries
ggplot() +
  geom_sf(data = disaggr_layer,
          fill = disaggr_layer$colors,
          color = disaggr_layer$colors) +
  geom_sf(data = county_layer,
          fill = NA,
          color = "#bbbbbb",
          size = .3) +
  geom_sf(data = state_layer,
          fill = NA) +
  theme_bw()
```

Most generally, we observe strong regional differences in mental health across the United States. Those differences range from an average of only 0.58 poor-mental-health days per month in Hand County, South Dakota, to 8.09 poor-mental-health days per month in Dickenson County, Virginia. This statement exemplifies the major advantage of the conventional mapping technique: We can directly infer the level of mental health for specific geographical units. For example, we can determine that Fulton County in Atlanta (Georgia's most populous county, which we frequently return to throughout this tutorial) falls into the category of 2.75 to 3.0 poor-mental-health days per month (exact score = 2.77). However, the current map is also limited in two crucial ways. First, in many (if not most) instances the aim of mapping is to reveal and communicate spatial patterns (in our case, the geographical distribution of mental health across the United States). Reporting scores separately on a county-by-county basis produces a scattered pattern that makes it difficult for readers to identify spatial regularities within the data. Second, the map does not cover the entire United States but contains blank spots for those counties that do not meet our minimum inclusion criteria. Thus, the current map does not provide any estimates on the mental-health levels of these areas.

counties. To balance these aims, for our data size and structure, a cutoff of 50 observations per county forms a reasonable compromise: Previous research has successfully employed this cutoff, which suggests that 50 observations per county are sufficiently high to inspire confidence in the resulting county-level mental-health scores. At the same time, employing a rather liberal cutoff of 50 observations per county allowed us to retain most counties ($N = 2,204$). Thus, we could at least partly avoid the problem that setting fixed cutoffs comes at the expense of selectively dropping observations with specific features from the analysis (i.e., usually less populous rural areas).

To map out differences in mental health, we broke up the county-level mental-health scores into equally sized categories and assigned darkening colors from yellow (low values) to red (high values). The following code maps out county-level mental-health scores for each county with at least 50 participants. Thereby, we use a color-coding scheme that groups counties in steps of 0.25 and assigns darkening colors from yellow to red.[3] Figure 1a depicts the resulting map.

Figure 1b is based on the same underlying data as the conventional map in Figure 1a. However, it is obvious that in this map, the data are visualized quite differently. In the following, we (a) explain distance-based weighting as the approach underlying this alternative visualization, (b) demonstrate how to design distance-based weighting so it is effective for geo-psychological research, and (c) discuss the advantages and disadvantages of this alternative visualization compared with the conventional mapping approach.

Mapping Based on Distance Weights

We now introduce distance-based weighting as an alternative technique to visualize spatial data. Distance-based weighting is a commonly used technique in geography and the wider social sciences (e.g., Buddemeier et al., 2011; Hough et al., 2004; Mann et al., 2002). Although an emerging set of recent studies highlighted the promise of distance-based weighting for mapping psychological data (Buecker et al., 2021; Calanchini et al., in press; Ebert et al., 2021, 2022), so far, this technique has been used in only a very few studies and by a narrow circle of researchers.

The basic idea of distance-based weighting is to calculate area estimates that represent distance-weighted averages of other measurement locations in the data. Thereby, following Tobler's (1970) first law of geography (i.e., "Everything is related to everything else. But near things are more related than distant things," p. 236), proximal locations receive greater weights than distal locations. Note that there is a great flexibility in how to exactly design this distance-based weighting. Below we outline in two steps how to design distance-based weighting so it is effective for geo-psychological research.

Step 1: transforming polygon data into point data on a spatial grid

The map presented in Figure 1 is based on the division of the contiguous United States in 3,108 counties. In our first analytical step, we now dissolve this division into counties and instead divide the surface of the United States into a regular spatial grid (i.e., into a raster of equally sized quadratic squares). Specifically, we use a fine-grained spatial grid dividing the United States into 57,570 squares with a size of 15 × 15 km (see thin dotted points in Fig. 2).[4] The following code creates this spatial raster:
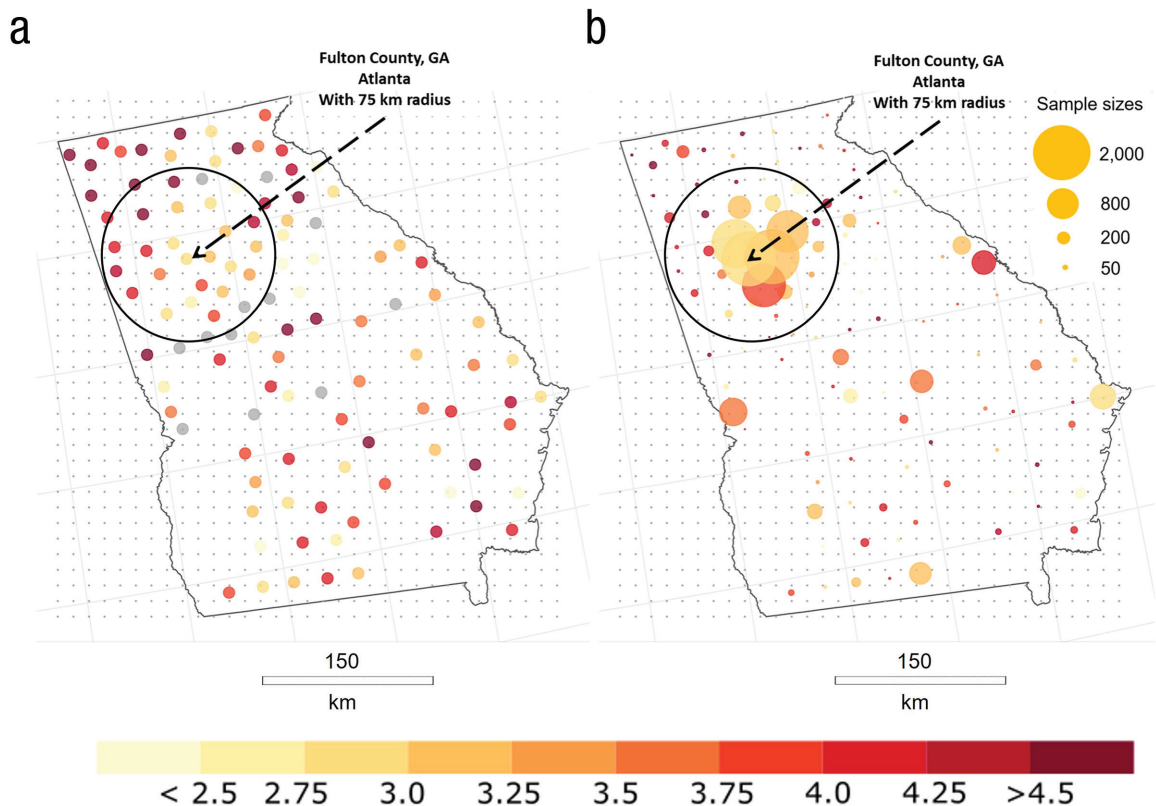

Fig. 2. (a) Transformation of the surface of Georgia into equally sized 15 × 15 km spatial grids with county scores superimposed as point data. Gray dots represent locations with insufficient sample sizes. (b) Spatial grid with measurement locations scaled according to their underlying sample sizes. In both panels, the black circle reflects a 75-km radius around the Fulton County, Georgia, measurement point location.

```
# Define grid size (in meters) and
  create raster layer
raster_size <- 15000
raster_layer <- raster(ext = extent
  (county_layer), res = c(raster_size,
  raster_size))
```

This spatial grid layer forms the basis for the second analytical step. Specifically, we superimpose the information depicted in Figure 1 (i.e., each county's average number of poor-mental-health days) onto this spatial grid.[5] Note that when doing so, we no longer treat county scores as polygon data but as point data. That is, we treat each county as a measurement point located in the geographical center of the specific county. Figure 2a displays how we conceptualize the data structure in the next analytical step. For example, we can see that Fulton County, Georgia, now represents one measurement location on a spatial grid. The gray dots represent those county scores that were excluded because their sample size did not meet our inclusion criterion of at least 50 participants.

Step 2: using distance weights to smooth point data

In our second analytical step, we now use distance-based weighting to smooth the data across the entire study area. To do so, we calculate a score for each grid cell, which represents a weighted average of all measurement points. Thereby, each measurement location is weighted according to its distance to that specific grid cell. To that end, we first need to calculate the distance between all grids and measurement points. The following code calculates these distances and stores them in a matrix:

```
# Create layer for distance-based
  weighting by removing counties
  without any observations
db_layer <- county %>%
  filter(!is.na(sp_var)) %>%
  as_Spatial()

# Define function to calculate
  Euclidean distances
eucl_dist <- function(p, q) {
  a <- outer(p[, 1], q[, 1], "-")^2
  b <- outer(p[, 2], q[, 2], "-")^2
  sqrt(a + b)
}

# Apply function to data points (here
  the centroids of the grids and the
  counties)
```

```
gdistmat <- eucl_dist(coordinates
  (raster_layer), coordinates(db_
  layer))
gdistmat <- gdistmat / 1000 # Transform
  distances to kilometers
```

Now that we know the distances between each grid cell and measurement location, we need to define how these distances translate into weights. In the simplest (and commonly used) approach of distance-based weighting, each location receives a weight, *w*, which represents the inverse of its distance to the target grid cell (also called inverse distance weighting),[6] which is given by:

$$w = \frac{1}{distance}.$$

The solid gray line in Figure 3 depicts the distance-decay function resulting from this weighting scheme, which shows a decline in assigned weights that trend toward zero. We can easily and flexibly adjust this distance-decay function. For example, by exponentiating distance (i.e., distance^2), the function would trend toward zero more quickly and therefore increase the relevance of nearby measurement points (see dashed gray line in Fig. 3).

To map out geo-psychological differences, we recommend adjusting this basic weighting scheme in two ways. First, we recommend specifying a distance-decay function that is adequate for the data structure and psychologically meaningful. Regarding adequacy, the used distance-decay function should be appropriate for the underlying data structure. For example, consider a case in which measurement locations are spaced (on average) 50 km apart from each other. In such a case, it would be suboptimal to use a distance-decay function that assigns high weights only to locations in very close distance (e.g., only to those within a 5-km radius). If we did so, single measurement locations will become overly dominant, whereas most others would not be relevant at all. Consequently, there would not be much of a smoothing effect. Thus, when informing a distance-decay function, we should try to ensure that most grids actually have measurement locations in the relevant spatial range. Regarding meaningfulness, the used distance-decay function should ideally be informed by theoretical considerations. For example, we here aim to map out regional differences across the United States. One approach to identify a person's psychologically relevant region is to define it as the interaction space that is available on a daily basis (Ebert et al., 2022; Götz et al., 2020; Stieger et al., 2022). Studies have shown that commuting or traveling for short-term activities is perceived as cumbersome if it exceeds a physical
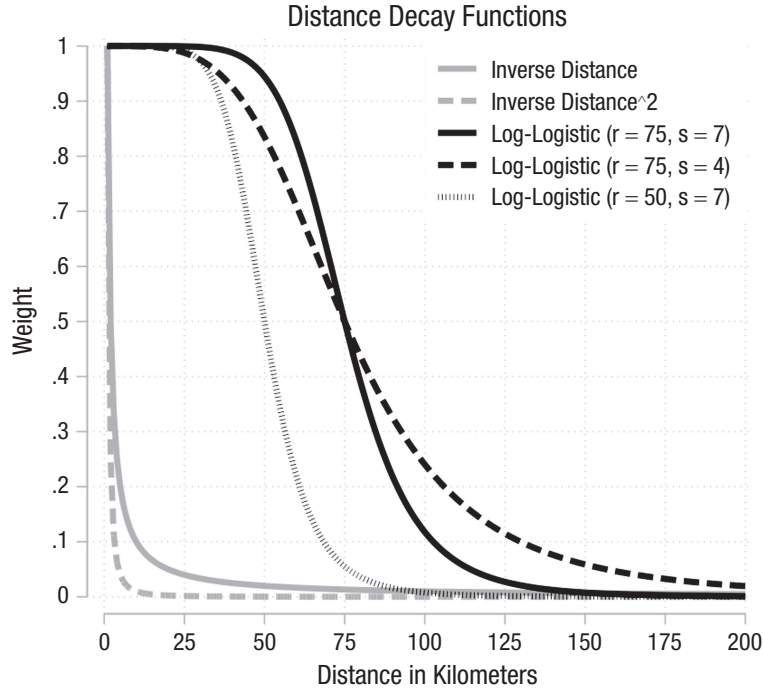
Distance Decay Functions

Fig. 3. Shapes of different distance-decay functions.

distance of 60 to 80 km (Ahmed & Stopher, 2014; Phibbs & Luft, 1995). A quick check shows that such a definition of distance should be adequate for our data structure. To illustrate, we can see in Figure 2 (75-km circle in Figs. 1a and 1b) that in most cases, several measurement points fall within the daily interaction space around a grid cell.

Note that there are many possibilities to define distance-decay functions (for an overview of commonly used approaches, see Brenner, 2017). In the following, we depict the daily available interaction space by using a log-logistic distance-decay function, which is given by

$$w = \frac{1}{1 + \left(\dfrac{distance}{r}\right)^s},$$

where r denotes the distance at which the decay function reaches a value of .5 and s determines the slope of the decay with distance. We use and introduce this distance-decay function here because it has the advantageous property that its general shape (flat – decreasing – flat) aligns neatly with the idea of representing the daily available interaction space. To illustrate, we defined the daily available interaction space to cover a distance of 60 to 80 km. If we chose $r = 75$ km and $s = 7$, this results in a distance-decay function in which (a) observations up to a distance of 50 km (i.e., within the core of inter-action space) receive a weight of nearly 1, (b) observations between 50 and 100 km (i.e., at the edges of the

interaction space) receive decreasing weights, and (c) observations beyond 100 km (i.e., clearly outside the interaction space) receive weights near 0. A further advantage of this log-logistic distance-decay function is that it can flexibly be adjusted to accommodate different research settings (see Fig. 3). For example, if we were interested in more fine-grained differences (e.g., within cities), we could set a lower r to shift the curve to the left and put more emphasis on nearby locations. Like-wise, we can adjust the shape of the distance function. For example, if we set $s = 4$ (instead of 7), the curve's trend toward 0 is less steep.[7]

Second, besides specifying an adequate and meaning-ful distance-decay function, we also recommend consid-ering the sample sizes of measurement locations (Brenner, 2017). Unlike, for instance, weather data (in which each location represents exactly the same number of measurements), in geo-psychological studies, loca-tions typically represent different numbers of partici-pants (and thus measurements). Consider Figure 2b, which shows that the measurement location of Fulton County represents many more participants than mea-surement locations from smaller rural communities in Georgia. To account for this, we may expand our weight-ing scheme by additionally weighting for sample size:

$$w = \frac{1}{1 + \left(\dfrac{distance}{75}\right)^7} \times samplesize$$

Including sample size in the equation comes with two advantages in our research context. First, we can now use all available data and must no longer exclude measurement locations with small sample sizes. Why is that? If we treat each measurement location equally (i.e., do not account for sample size), small samples with unreliable means could severely bias the estimation (which is why they are excluded). However, this is no longer the case if we additionally weigh for sample size. Specifically, measurement locations with very few people will be less influential because they receive small-sample-size weights (see Fig. 2: gray dots in Fig. 2a are transformed to small colored dots in Fig. 2b). Second, weighting according to sample size allows for an easy and straightforward interpretation of the resulting map: For each grid cell, we now indicate the average score among the people in our data that reside within the daily available interaction space of that grid cell. The following code defines the weighting function and uses the distance matrix (created in the previous step) to calculate a weighted score for each grid cell:

```
# Define function to calculate spatial
  weights based on log logistic
  distance decay function
spatial_weights <- function(d, r, s) {
  (1/(1+((d/r)^s)))
}
r <- 75
s <- 7

# Compute spatial weights based on the
  computed geographic distances between
  grid cells and observed counties
gweights <- spatial_weights(gdistmat,
  r = r, s = s)

# Define function that takes spatial
  weights (w), the desired variable
  (v), and the number of participants
  (n) as input parameters
db_fun <- function(w, v, n) {
  x <- rowSums(t(apply(w, 1,
    function(x) x * v * n)), na.rm =
    TRUE)
  x_weight <- rowSums(t(apply(w, 1,
    function(x) x * n)), na.rm = TRUE)
  x <- x / x_weight
  x
}

# Apply function to compute distance-
  based weighted scores for every grid
  cell
```

```
db_values <- db_fun(w = gweights, v =
  db_layer$sp_var, n = db_layer$nobs)

# Assign distance-based weighted
  scores to grid cells
db_raster_layer <- setValues(raster_
  layer, db_values)
```

Mapping differences in mental health using spatial grids and distance weighting

Because of the previous step, we now have a weighted average score for each grid cell and are ready for mapping. The following code prepares the spatial grid and maps out grid-cell scores using the same grouping and color scheme that we used for the conventional map (i.e., groups of 0.25 with darkening colors from yellow to red). Figure 1b depicts the resulting map and juxtaposes it against the conventional mapping approach (Fig. 1a):

```
# Crop raster layer to the shape of
  the contiguous US
db_raster_layer <- mask(db_raster_
  layer, county)

# Transform raster layer to data frame
db_raster_layer <- as.data.frame
  (db_raster_layer, xy = TRUE)
db_raster_layer <- na.omit
  (db_raster_layer)

# Group calculated values using the
  color coding scheme from above
db_raster_layer$group <- findInterval
  (db_raster_layer$layer,
  grouping) + 1
db_raster_layer$colors <- colors
  [db_raster_layer$group]

# Create map
ggplot() +
  geom_raster(data = db_raster_layer,
    aes(x = x, y = y), fill = db_raster_
    layer$colors, color = db_raster_
    layer$colors) +
  geom_sf(data = state_layer,
          fill = NA) +
  labs(x = "",
       y = "") +
  theme_bw()
```

When comparing distance-based weighting with the conventionally used mapping technique (Fig. 1a vs.

Table 1. Comparison of Mapping Techniques

Conventional mapping	Distance-based weighting
• Easy to implement	• Easy to implement
• Easy to interpret	• Easy to interpret
• Provides scores for specific geographical units	• Highlights distributional patterns
	• Weighting function captures psychological contexts
	• Not tied to a priori specified units
	• Estimates for areas with sparse data

Fig. 1b), we first realize that it is no longer possible to infer exact scores for specific geographical units (e.g., What is the average score of the participants residing in Fulton County?). Any prespecified borders of geographical units are dissolved. Furthermore, because of the distance-based weighting, the score for any specific location is now also informed by nearby locations.[8] This smoothing generally leads to less extreme spatial differences. Specifically, whereas scores in Figure 1a range from 0.58 to 8.09 (Δ = 7.51), in Figure 1b, they here range from 1.90 to 6.30 (Δ = 4.4).

Although it is no longer possible to infer exact scores for specific locations, the map in Figure 1b provides clear advantages in communicating general spatial patterns in the data. Specifically, at least three merits stand out. First, distance-based weighting leads to continuous, easy-to-grasp distributional patterns. Thus, distributional patterns that were previously disguised (or only tentatively visible) now become evident at a single glance. To illustrate, the map clearly indicates a belt of poor mental health spanning from Oklahoma through Missouri and Arkansas, into the Deep South, and back up to Kentucky and West Virginia. Second, according to the very fine-grained spatial grid, the depicted distributional patterns are no longer tied to prefixed county boundaries. Rather, the distributional patterns could largely emerge freely from the data using our weighting scheme. This weighting scheme was informed by theoretical considerations. Thus, the distributional patterns should reflect "true" psychological contexts. To illustrate, the map clearly depicts how the extensive belt of poor mental health in the South is broken up by the sphere of influence of agglomerations such as Atlanta. Third, because we calculated scores for each grid cell, we now have estimates for areas in which only sparse data are available (i.e., the areas that were blank in Fig. 1a). To illustrate, consider the extensive areas of Texas that were missing data in Figure 1a and now form a nuanced geographical distribution.[9]

Taken together, we have shown that both mapping techniques (i.e., conventional mapping and distance-based weighting) come with their own unique strengths and limitations. Table 1 summarizes the major strengths of both approaches. In short, if researchers are most interested in communicating precise estimates for specific geographical units, color-coding disaggregated data will serve them well. Conversely, if the aim is to communicate broader patterns, distance-based weighting may be preferred. Likewise, researchers may also complement different techniques by presenting their results side by side. In sum, we neither discourage the use of conventional mapping nor advocate for a universal use of distance-based weighting. Which mapping technique is most appropriate will ultimately depend on the research question at hand.

Conclusion

Psychologists of all subfields are becoming increasingly interested in the geographical distribution of psychological phenomena. To visualize such geo-psychological differences, previous studies largely relied on conventional mapping (i.e., color-coding disaggregated data). In this tutorial, we reviewed the strengths and limitations of this basic mapping technique (i.e., exact scores for specific locations but spatial regularities and blank areas that are difficult to identify). We then introduced distance-based weighting as an alternative mapping technique that not many psychologists are aware of yet. In two steps, we demonstrated how to effectively employ this technique in geo-psychological research. First, we transformed polygon data into point data on a spatial grid. This transformation allowed us to produce more continuous maps with distributions that are not tied to prefixed higher spatial entities. Second, we used distance weights to estimate scores for each spatial grid cell. This allowed us to (a) smooth the data (i.e., making it easier to identify spatial regularities) and (b) estimate scores for unmeasured locations (i.e., filling in blank areas in the map).

From our perspective, distance-based weighting can be a useful tool for geo-psychological research for at least three reasons. First, distance-based weighting is very versatile and can be applied regardless of which construct is being studied (e.g., personality, bias, or well-being), how it was assessed (e.g., count variables, binary data, or Likert scales), and at what geographical level it is being examined (e.g., from neighborhood differences

in cities to regional differences in countries). Second, as demonstrated by our code, distance-based weighting is relatively easy to implement. Third (and potentially most relevant), distance-based weighting is straightforward to interpret. For example, in our case, we indicate for each place the average score among the people in the data that reside within the daily available interaction space. We consider this straightforward interpretation particularly advantageous when the audience of an article is nonexperts in spatial modeling (which is almost always the case in geo-psychological research). That being said, there naturally exist further approaches to visualize and interpolate spatial data. Box 1 provides an annotated list of five tools that may be useful for researchers studying spatial differences in psychological phenomena. All these tools are implemented in major software packages (e.g., R) and can be applied without additional costs (see e.g., Bivand et al., 2013; Lopez-Martin et al., 2021).

Box 1. List of Further Useful Tools for Geo-Psychological Research

---

Hot-Spot Analysis

The aim of hot-spot analysis is to visualize and statistically identify spatial clusters (i.e., areas in which high/low values of a variable flock together). Hot-spot analysis indicates for each geographical unit the degree to which this unit and its nearby units deviate from the overall study area. This approach has already been used in geo-psychological research (e.g., Ebert et al., 2019; Jokela et al., 2015; Rentfrow et al., 2015). For detailed information on this method, see Bivand and Wong (2018).

Nearest Neighbor Interpolation (or Proximal Interpolation)

Nearest neighbor interpolation is the most basic and oldest form of spatial interpolation. It can be understood as a special case of distance-based weighting: Every spot on the map simply receives the value of its nearest measurement location. As such, nearest neighbor interpolation uses only little of the available information, but it is easy to implement and interpret. For example, we could easily implement nearest neighbor interpolation in our case by assigning each grid cell the value of the nearest measurement location. In Supplemental Script 5, we provide code demonstrating how to implement this basic method.

Trend-Surface Analysis

The idea of trend-surface analysis is to interpolate data according to the principle that the distribution of a variable changes along a geographic gradient. To model this, the variable of interest is entered in a regression model as the dependent variable and latitude and longitude are the independent variables. By specifying higher-order polynomials, also complex distributional patterns can be captured. To our knowledge, trend-surface analysis has not yet been used in geo-psychological research. For detailed information on this method, see Agterberg (1984).

Kriging (or Gaussian Process Regression)

Kriging also uses information from nearby geographical units to interpolate the value of a variable at a place. However, unlike the inverse distance weighting presented in the current tutorial, Kriging is a nondeterministic approach of spatial interpolation. That is, the distance-based weights are not defined a priori but are fitted to the spatial autocorrelation structure in the data using so-called variograms. Because of this more complex modeling, Kriging minimizes the estimation variance at each estimated location. There are different forms of Kriging, and its application requires advanced geo-statistical expertise. To our knowledge, Kriging has not yet been applied in geo-psychological research. For detailed information on Kriging, see Oliver and Webster (2015).

Multilevel Regression and (Synthetic) Poststratification

Multilevel regression and poststratification (MrP; Gelman & Little, 1997) and its more recent extension, multilevel regression with synthetic poststratification (MrsP; Leemann & Wasserfallen, 2017), are quickly gaining momentum as sophisticated tools for spatial imputation. Researchers derive predictions for an outcome variable in specific sociodemographic brackets (e.g., 40- to 65-year-old men with a university degree). Drawing from granular census data, researchers then compute area estimates in the outcome variable by weighting the prediction of each sociodemographic bracket by the number of residents in an area that fall into the corresponding bracket. This process originally stems from political science, and Hoover and Dehghani (2020) provided an excellent introduction for the use of Mr(s)P in geo-psychological research.

Transparency

ORCID iD

Tobias Ebert https://orcid.org/0000-0003-0146-2517

Notes

1. Note that throughout the current article, the term "county" is used to capture both counties and county equivalents. County equivalents refer to areas that are comparable with counties but are officially called by other names (e.g., parishes in Louisiana, councils of government in Connecticut, or independent cities in Maryland, Missouri, Nevada, and Virginia; Pew Research Center, 2022; United States Department of Commerce, 1994).

2. We also provide data sets containing the additional variables *smoker* (county average from binary variable 0 = nonsmoker, 1 = smoker) and *life_satis* (In general, how satisfied are you with your life? county-average from Likert scale variable 1 [*very satisfied*] to 4 [*very dissatisfied*]). With these data, readers can reproduce our analytical steps and explore maps for variables with different content and scaling.

3. Note that there are many different options on how to group and color-code a measurement across geographical areas. Which categorization and color scheme to use will ultimately depend on the distribution of the underlying variable and the communication purpose (e.g., researchers that want to depict general distributional patterns may choose a different categorization and color scheme than researchers that want to highlight where the most extreme scores are located).

4. In the subsequent analytical step, we will calculate the distance between each grid cell and measurement location pairing.

Thus, the computational intensity grows exponentially with the number of grid cells. Therefore, when defining their spatial raster, researchers should keep in mind that choosing an unnecessarily fine-grained spatial raster may not be computationally feasible. However, in most research settings that we can think of, computational intensity should not be a serious bottleneck. Even our subsequent analysis (which is based on the fairly high number of 57,570 grid cells and 2,360 measurement locations) should run smoothly on most conventional machines. Nevertheless, there might be situations in which the underlying data structure (i.e., the combination of grid cells and measurement locations) exceeds computational feasibility. In Supplemental Script 6 we provide a workaround for such situations by excluding those measurement locations that are not informative for a grid cell's score.

5. Note that it is not strictly necessary to dissolve the county polygons into spatial grids first. It would also be possible to superimpose the county scores as point data onto the county map. However, using a spatial grid leads to more fine-grained, continuous distributional patterns that are not tied to prefixed spatial boundaries.

6. Physical distance is by far the most commonly used indicator to capture the proximity between two locations. However, in some cases (e.g., when two places are located at opposite sides of a river), the geographical distance may not well capture the "psychological" distance between two locations (Murdani et al., 2018). Therefore, note that it is also possible to use other definitions of distance, such as the travel time between two locations (e.g., Duschl et al., 2015; Mewes & Ebert, 2021). However, at least to our knowledge, there currently exists no other option to extract travel time other than repeatedly calling public routing application programming interfaces from, for example, Google Maps or Open Street Map. Accordingly, calculating the distance between hundreds of thousands of location pairings takes very long and frequently exceeds the maximum number of allowed requests of the public routing services. Whether it is worth going to these lengths ultimately depends on the research question at hand. Note that often it will not make much of a difference whether the visualization is based on geographical distance or travel time. That is because travel time is often a linear function of geographical distance. To illustrate, Ebert et al. (2022) reported that the correlation between travel time and geographical distance across a randomly drawn set of U.S. zip code areas was $r = .99$ and across zip codes within the New York area was $r = .92$.

7. In Supplemental Script 3, we provide additional maps to show how distributional patterns become more fine-grained and scattered when choosing a smaller r and/or a higher s and more continuous and smoother when choosing a higher r and/or a smaller s.

8. Naturally, some grid cells are backed up by more nearby measurement locations than others. Thus, we may still seek to exclude grid cells that are informed only by very sparse nearby data. To do so, we can blank all grid cells whose cumulated weight sum of participants is below a certain threshold. In Supplemental Script 4, we provide additional maps in which we excluded grid cells whose cumulated weight sum was below 50 and 100, respectively.

9. Distance-based weighting can also be described as a simple, deterministic approach of data interpolation. In Supplemental

Script 2, we provide code to demonstrate how to extract grid-cell scores for points (i.e., what the grid-cell score is at a specific point in space) and polygons (i.e., what the average grid-cell score is within a specific bounding box such as a zip code area or a county) so they can be used in further analytical steps. However, we note that we here introduced distance-based weighting as an approach to visualize data. When the aim is to interpolate data, other more complex approaches may be more powerful (see Box 1).

References

Agterberg, F. P. (1984). Trend surface analysis. In G. L. Gaile & C. J. Willmott (Eds.), *Spatial statistics and models* (pp. 147–171). Springer.

Ahmed, A., & Stopher, P. (2014). Seventy minutes plus or minus 10 – A review of travel time budget studies. *Transport Reviews*, *34*(5), 607–625.

Bivand, R. S., Pebesma, E. J., & Gómez-Rubio, V. (2013). *Applied spatial data analysis with R*. Springer.

Bivand, R. S., & Wong, D. W. (2018). Comparing implementations of global and local indicators of spatial association. *Test*, *27*, 716–748.

Bleidorn, W., Schönbrodt, F., Gebauer, J. E., Rentfrow, P. J., Potter, J., & Gosling, S. D. (2016). To live among like-minded others: Exploring the links between person–city personality fit and self-esteem. *Psychological Science*, *27*(3), 419–427.

Brenner, T. (2017). *Identification of clusters – An actor-based approach* (Working Papers on Innovation and Space, Vol. 02.17) (forthcoming in Cantner, U., & Guerzoni, M. (Eds.). *Handbook of research methods and applications in industrial dynamics and evolutionary economics*. Edward Elgar).

Buddemeier, R. W., Lane, D. R., & Martinich, J. A. (2011). Modeling regional coral reef responses to global warming and changes in ocean chemistry: Caribbean case study. *Climatic Change*, *109*(3), 375–397.

Buecker, S., Ebert, T., Götz, F. M., Entringer, T. M., & Luhmann, M. (2021). In a lonely place: Investigating regional differences in loneliness. *Social Psychological and Personality Science*, *12*, 147–155.

Calanchini, J., Hehman, E., Ebert, T., Wilson, L., Simon, D., & Esposito, E. (in press). Regional intergroup bias. *Advances in Experimental Social Psychology*. https://psyarxiv.com/gh9vj/download?format=pdf

Centers for Disease Control and Prevention. (2020, January 4). *Behavioral risk factor surveillance system*. https://www.cdc.gov/brfss/index.html

Chopik, W. J., & Motyl, M. (2017). Is Virginia for lovers? Geographic variation in adult attachment orientation. *Journal of Research in Personality*, *66*, 38–45.

Duschl, M., Scholl, T., Brenner, T., Luxen, D., & Raschke, F. (2015). Industry-specific firm growth and agglomeration. *Regional Studies*, *49*, 1822–1839.

Ebert, T., Gebauer, J. E., Brenner, T., Bleidorn, W., Gosling, S. D., Potter, J., & Rentfrow, P. J. (2022). Are regional differences in personality and their correlates robust? Applying spatial analysis techniques to examine regional variation in personality. *Perspectives on Psychological Science*, *17*(2), 407–441. https://doi.org/10.1177/1745691621998326

Ebert, T., Götz, F. M., Gladstone, J. J., Müller, S. R., & Matz, S. C. (2021). Spending reflects not only who we are but also who we are around: The joint effects of individual and geographic personality on consumption. *Journal of Personality and Social Psychology*, *121*(2), 378–393. https://doi.org/10.1037/pspp0000344

Ebert, T., Götz, F. M., Mewes, L., & Rentfrow, P. J. (2022). Spatial analysis for psychologists: How to use individual-level data for research at the geographically aggregated level. *Psychological Methods*. Advance online publication. https://doi.org/10.1037/met0000493

Ebert, T., Götz, F. M., Obschonka, M., Zmigrod, L., & Rentfrow, P. J. (2019). Regional variation in courage and entrepreneurship: The contrasting role of courage for the emergence and survival of start-ups in the United States. *Journal of Personality*, *87*(5), 1039–1055.

Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., Jha, S., Agrawal, M., Dziurzynski, L. A., Sap, M., Weeg, C., Larson, E. E., Ungar, L. H., & Seligman, M. E. P. (2015). Psychological language on twitter predicts county-level heart disease mortality. *Psychological Science*, *26*(2), 159–169.

Elleman, L. G., Condon, D. M., Holtzman, N. S., Allen, V. R., & Revelle, W. (2020). Smaller is better: Associations between personality and demographics are improved by examining narrower traits and regions. *Collabra: Psychology*, *6*(1), Article 17210. https://doi.org/10.1525/collabra.17210

Gelman, A., & Little, T. C. (1997). Poststratification into many categories using hierarchical logistic regression. *Survey Research*, *23*, 127–135.

Götz, F. M., Ebert, T., & Rentfrow, P. J. (2018). Regional cultures and the psychological geography of Switzerland: Person-environment-fit in personality predicts subjective wellbeing. *Frontiers in Psychology*, *9*, Article 517. https://doi.org/10.3389/fpsyg.2018.00517

Götz, F. M., Stieger, S., Gosling, S. D., Potter, J., & Rentfrow, P. J. (2020). Physical topography is associated with human personality. *Nature Human Behaviour*, *4*, 1135–1144.

Hehman, E., Calanchini, J., Flake, J. K., & Leitner, J. B. (2019). Establishing construct validity evidence for regional measures of explicit and implicit racial bias. *Journal of Experimental Psychology: General*, *148*(6), 1022–1040.

Hehman, E., Flake, J. K., & Calanchini, J. (2017). Disproportionate use of lethal force in policing is associated with regional racial biases of residents. *Social Psychological and Personality Science*, *9*(4), 393–401.

Hehman, E., Ofosu, E. K., & Calanchini, J. (2020). Using environmental features to maximize prediction of regional intergroup bias. *Social Psychological and Personality Science*, *12*(2), 156–164.

Hoover, J., & Dehghani, M. (2020). The big, the bad, and the ugly: Geographic estimation with flawed psychological data. *Psychological Methods*, *25*, 412–419.

Hough, R. L., Breward, N., Young, S. D., Crout, N. M., Tye, A. M., Moir, A. M., & Thornton, I. (2004). Assessing potential risk of heavy metal exposure from consumption of

home-produced vegetables by urban populations. *Environmental Health Perspectives*, *112*(2), 215–221.

Jaidka, K., Giorgi, S., Schwartz, H. A., Kern, M. L., Ungar, L. H., & Eichstaedt, J. C. (2020). Estimating geographic subjective well-being from Twitter: A comparison of dictionary and data-driven language methods. *Proceedings of the National Academy of Sciences, USA*, *117*(19), 10165–10171.

Jokela, M., Bleidorn, W., Lamb, M. E., Gosling, S. D., & Rentfrow, P. J. (2015). Geographically varying associations between personality and life satisfaction in the London metropolitan area. *Proceedings of the National Academy of Sciences, USA*, *112*, 725–730.

Leemann, L., & Wasserfallen, F. (2017). Extending the use and prediction precision of subnational public opinion estimation. *American Journal of Political Science*, *61*, 1003–1022.

Lopez-Martin, J., Phillips, J. H., & Gelman, A. (2021). *Multilevel regression and poststratification case studies*. https://book down.org/jl5522/MRP-case-studies/

Mann, J. K., Tager, I. B., Lurmann, F., Segal, M., Quesenberry, C. P., Jr., Lugg, M. M., Shan, J., & Van Den Eeden, S. K. (2002). Air pollution and hospital admissions for ischemic heart disease in persons with congestive heart failure or arrhythmia. *Environmental Health Perspectives*, *110*(12), 1247–1252.

Matz, S. C., & Gladstone, J. J. (2020). Nice guys finish last: When and why agreeableness is associated with economic hardship. *Journal of Personality and Social Psychology*, *118*(3), 545–561.

Mewes, L., & Ebert, T. (2021). Where do people want to become entrepreneurs? Mapping fine-grained geographical differences in entrepreneurial potential across Great Britain. *Regional Studies, Regional Science*, *8*, 332–335.

Murdani, M. H., Kwon, J., Choi, Y. H., & Hong, B. (2018). Efficient proximity computation techniques using zip code data for smart cities. *Sensors*, *18*(4), Article 965. https://doi.org/10.3390/s18040965

Obschonka, M., Lee, N., Rodríguez-Pose, A., Eichstaedt, J. C., & Ebert, T. (2019). Big data methods, social media, and the psychology of entrepreneurial regions: Capturing cross-county personality traits and their impact on entrepreneurship in the USA. *Small Business Economics*, *55*(3), 567–588.

Obschonka, M., Stuetzer, M., Rentfrow, P. J., Lee, N., Potter, J., & Gosling, S. D. (2018). Fear, populism, and the geopolitical landscape: The "sleeper effect" of neurotic personality traits on regional voting behavior in the 2016 Brexit and Trump elections. *Social Psychological and Personality Science*, *9*(3), 285–298.

Ofosu, E. K., Chambers, M. K., Chen, J. M., & Hehman, E. (2019). Same-sex marriage legalization associated with reduced implicit and explicit antigay bias. *Proceedings of the National Academy of Sciences, USA*, *116*, 8846–8851.

Oishi, S., & Graham, J. (2010). Social ecology: Lost and found in psychological science. *Perspectives on Psychological Science*, *5*(4), 356–377.

Oliver, M. A., & Webster, R. (2015). *Basic steps in geostatistics: The Variogram and Kriging*. Springer International Publishing.

Payne, B. K., Vuletich, H. A., & Brown-Iannuzzi, J. L. (2019). Historical roots of implicit bias in slavery. *Proceedings of the National Academy of Sciences, USA*, *116*(24), 11693–11698.

Pew Research Center. (2022, February 17). *County definitions for 1980–2020*. https://www.pewresearch.org/hispanic/wp-content/uploads/sites/5/2022/01/RE_2022.01.31_Hispanic-population-County-definitions_FINAL.pdf

Phibbs, C., & Luft, H. (1995). Correlation of travel time on roads versus straight line distance. *Medical Care Research and Review*, *52*, 532–542.

Rentfrow, P. J. (2020). Geographical psychology. *Current Opinion in Psychology*, *32*, 165–170.

Rentfrow, P. J., Gosling, S. D., Jokela, M., Stillwell, D. J., Kosinski, M., & Potter, J. (2013). Divided we stand: Three psychological regions of the United States and their political, economic, social, and health correlates. *Journal of Personality and Social Psychology*, *105*(6), 996–1012.

Rentfrow, P. J., Gosling, S. D., & Potter, J. (2008). A theory of the emergence, persistence, and expression of geographic variation in psychological characteristics. *Perspectives on Psychological Science*, *3*(5), 339–369.

Rentfrow, P. J., & Jokela, M. (2016). Geographical psychology: The spatial organization of psychological phenomena. *Current Directions in Psychological Science*, *25*(6), 393–398.

Rentfrow, P. J., Jokela, M., & Lamb, M. E. (2015). Regional personality differences in Great Britain. *PLOS ONE*, *10*(3), Article e0122245. https://doi.org/10.1371/journal.pone.0122245

Stieger, S., Götz, F. M., Wilson, C., Volsa, S., & Rentfrow, P. J. (2022). A tale of peaks and valleys: Sinusoid relationship patterns between mountainousness and basic human values. *Social Psychological and Personality Science*, *13*(2), 390–402. Advance online publication. https://doi.org/10.1177/19485506211034966

Tobler, W. (1970). A computer movie simulating urban growth in the Detroit region. *Economic Geography*, *46*, 234–240.

United States Census Bureau. (2020, September 14). *Geography program*. https://www.census.gov/programs-surveys/geography.html

United States Department of Commerce. (1994). *Geographic areas reference manual*. United States Department of Commerce, Economics and Statistics Administration, Bureau of the Census.