Axioms of a Metric

Picture analysis always assumes that pictures are defined in coordinates, and we apply the Euclidean metric as the "golden standard" for distance (or derived, such as area) measurements. However, grids may create difficulties sometimes for applying the Euclidean metric, or may suggest to use alternative metrics.

Let S be an arbitrary nonempty set. A function $d : S \times S \mapsto \mathbb{R}$ is a *distance function* or *metric* on S iff it has the following properties:

M1: For all $p, q \in S$, we have $d(p, q) \geq 0$, and $d(p, q) = 0$ iff $p = q$ (positive definiteness).

M2: For all $p, q \in S$, we have $d(p, q) = d(q, p)$ (symmetry).

M3: For all $p, q, r \in S$, we have $d(p, r) \leq d(p, q) + d(q, r)$ (triangularity: the triangle inequality).

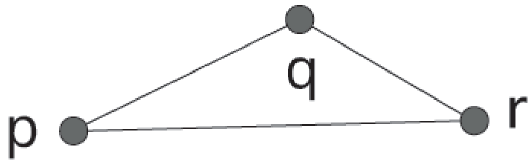

If d is a metric on S, the pair $[S, d]$ defines a *metric space*.

Let $[S, d]$ be a metric space, $p \in S$, and $\epsilon > 0$.

The set of points $q \in S$ such that $d(p, q) \leq \varepsilon$ is called a *ball* of radius ε with *center* p (or a *disk* if S is planar), and

$$U_\varepsilon(p) = \{q : \; q \in S \; \wedge \; d(p, q) < \varepsilon\}$$

is the ε-*neighborhood* of p in S.

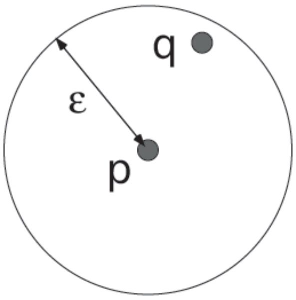

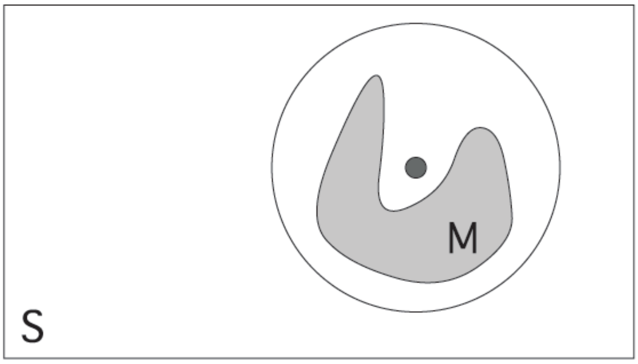

A subset M of S is called *bounded* iff it is contained in a ball of some finite radius.

Euclidean Metric

Euclidean space $\mathbb{E}^n$ is defined with an orthogonal coordinate system; it is used to define a metric d_e called the *Euclidean metric*:

$$d_e(p, q) = \sqrt{(x_1 - y_1)^2 + \ldots + (x_n - y_n)^2}$$

where points p, q have coordinates $(x_1, \ldots, x_n)$ and $(y_1, \ldots, y_n)$.

Theorem 1 *d_e is a metric on $\mathbb{R}^n$.*

A Degenerate Example

Let S be a nonempty set, and we define that

$$d_b\left(p,q\right) = \begin{cases} 0 & \text{if } p = q \\ 1 & \text{otherwise} \end{cases}$$

$[S, d_b]$ is a metric space (d_b is the *binary metric*):

M1:

M2:

M3:

$[S, d_b]$ is bounded:

Bounded Metric Spaces

$[\mathbb{R}^n, d_e]$ are unbounded metric spaces.

If $[S, d]$ is an unbounded metric space,

$$d'\left(p,q\right) = \frac{d(p,q)}{1 + d(p,q)}$$

defines a metric d' on S, and $[S, d']$ is a bounded metric space.

Integer-Valued Metrics

Assume a real a:

$$\begin{aligned}
&\lfloor a \rfloor, \text{ the largest integer less than or equal to } a \\
&\lceil a \rceil, \text{ the smallest integer greater than or equal to } a \\
&[a], \text{ the nearest integer to } a \text{ if it is unique,} \\
&\qquad\qquad \text{and } \lfloor a \rfloor \text{ otherwise}
\end{aligned}$$

For any function $d : S \times S \mapsto \mathbb{R}$, we can define $\lfloor d \rfloor$ by $\lfloor d \rfloor(p, q) = \lfloor d(p, q) \rfloor$ and similarly for $\lceil d \rceil$ and $[d]$.

Even if d is a metric, these integer-valued functions may not be metrics.

Example: $\lfloor d_e \rfloor$ and $[d_e]$ are not metrics on $\mathbb{Z}^2$.

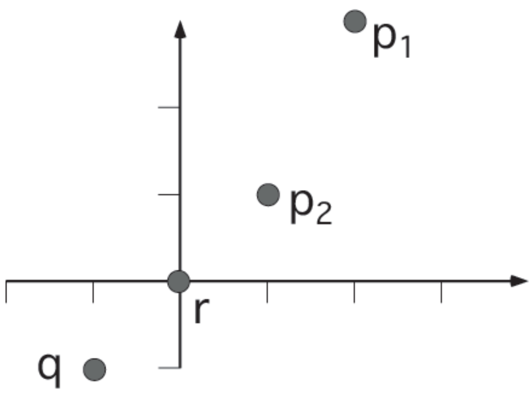

Theorem 2 *If d is a metric* $\lceil d \rceil$ *is also a metric.*

Minkowski Metrics

Minkowski metrics L_m on $\mathbb{R}^n$, where $L_2 = d_e$ (Euclidean metric):

$$L_m(p,q) = \sqrt[m]{|x_1 - y_1|^m + \ldots + |x_n - y_n|^m} \qquad (m = 1, 2, \ldots)$$
$$L_\infty(p,q) = \max\{|x_1 - y_1|, \ldots, |x_n - y_n|\}$$

where $p = (x_1, x_2, \ldots, x_n)$ and $q = (y_1, y_2, \ldots, y_n)$.

Theorem 3 *All the* L_m *are metrics on* $\mathbb{R}^n$*, with*

$$L_{m_1}(p,q) \leq L_{m_2}(p,q) \quad \textit{for all} \;\; 1 \leq m_2 \leq m_1 \leq \infty$$

and all $p, q \in \mathbb{R}^n$.

Relation to Grid Adjacencies

$L_1(p,q) = |x_1 - y_1| + \ldots + |x_n - y_n|$

Two 2D grid points p_1 and p_2 are 4-adjacent iff $L_1(p_1,p_2) = 1$ and 8-adjacent iff $L_\infty(p_1,p_2) = 1$.

Two 3D grid points p_1 and p_2 are 6-adjacent iff $L_1(p_1,p_2) = 1$ and 26-adjacent iff $L_\infty(p_1,p_2) = 1$.

2D Grid Point Metrics

Let $p, q \in \mathbb{R}^2$, $p = (x_1, y_1)$, $q = (x_2, y_2)$, and

$$d_4(p,q) = |x_1 - x_2| + |y_1 - y_2|$$

$[\mathbb{R}^2, d_4]$ is a metric space; d_4 is the Minkowski metric L_1.

We call d_4 the *city-block metric* or *Manhattan metric* because, when we restrict it to $\mathbb{Z}^2$, $d_4(p,q)$ is the minimal number of isothetic unit-length steps from p to q; it resembles a shortest walk in a city with streets that are laid out in an orthogonal grid pattern.

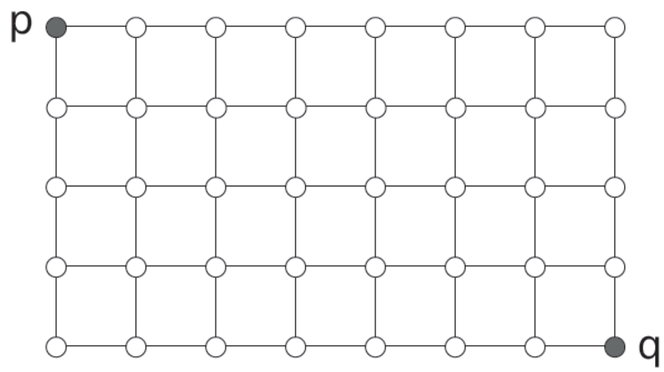

Let $p, q \in \mathbb{R}^2$, $p = (x_1, y_1)$, and $q = (x_2, y_2)$, and

$$d_8(p,q) = \max\{|x_1 - x_2|, |y_1 - y_2|\}$$

$[\mathbb{R}^2, d_8]$ is a metric space; d_8 is the Minkowski metric L_∞.

We call d_8 the *chessboard metric* because, when we restrict it to $\mathbb{Z}^2$, $d_8(p,q)$ is the minimal number of moves from p to q by a king on a chessboard.

Theorem 4 $d_8(p,q) \leq d_e(p,q) \leq d_4(p,q) \leq 2 \cdot d_8(p,q)$ for all $p, q \in \mathbb{R}^2$

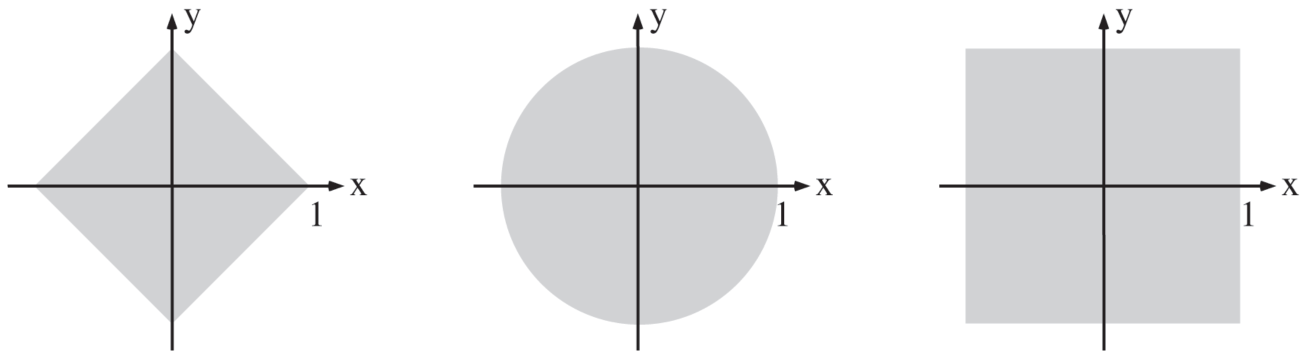

The city block, Euclidean, and chessboard unit disks in the real plane.

We have $U_\varepsilon(p) = \{p\}$ for $\varepsilon \leq 1$ and any of the metrics $\lceil d_e \rceil, d_4$ and d_8 on $\mathbb{Z}^2$, or for the binary metric (on any set). Metrics on grid cells are defined by identifying cells with their centers.

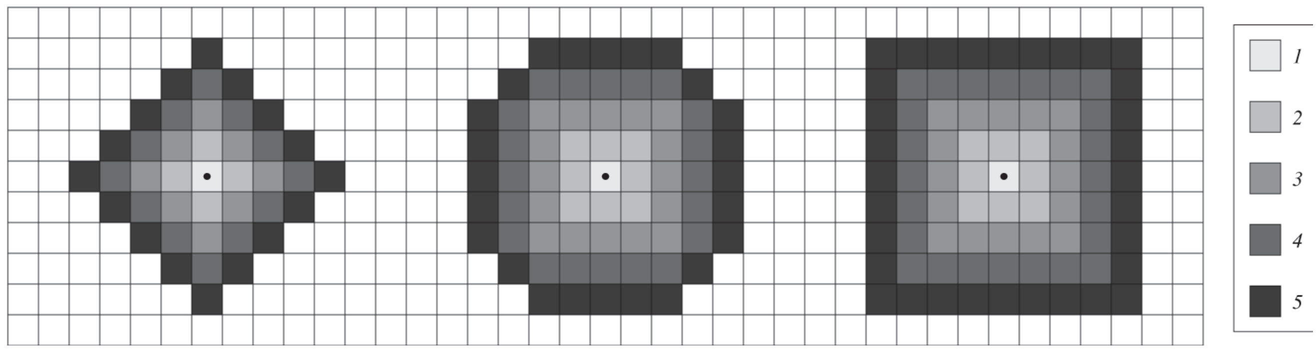

The ε-neighborhoods for $\varepsilon = 1, 2, 3, 4$, and 5 in the 2D grid cell model defined by the city block (left), Euclidean (middle), and chessboard (right) metrics.

For any grid point p, the smallest neighborhood of p in $[\mathbb{Z}^2, d_\alpha]$ ($\alpha \in \{4, 8\}$) is defined by

$$N_\alpha(p) = \{q \in \mathbb{Z}^2 : \ d_\alpha(p, q) \leq 1\}$$

3D Grid Point Metrics

Let $p, q \in \mathbb{R}^3, p = (x_1, y_1, z_1)$, and $q = (x_2, y_2, z_2)$, and

$$d_6(p, q) = L_1(p, q) = |x_1 - x_2| + |y_1 - y_2| + |z_1 - z_2|$$

$$d_{26}(p, q) = L_\infty(p, q) = \max\{|x_1 - x_2|, |y_1 - y_2|, |z_1 - z_2|\}$$

We also define

$$d_{18}(p, q) = \max\{d_{26}(p, q), \lceil d_6(p, q)/2 \rceil\}$$

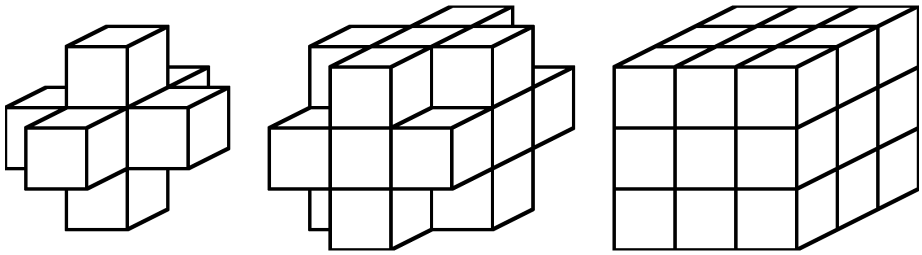

For any grid point p, the smallest neighborhood of p in $[\mathbb{Z}^3, d_\alpha]$ ($\alpha \in \{6, 18, 26\}$) is defined by

$$N_\alpha(p) = \{q \in \mathbb{Z}^3 : \ d_\alpha(p, q) \leq 1\}$$

$N_\alpha(p) - \{p\}$ has cardinality α for $\alpha = 6, 18$, and 26.

Theorem 5 $d_{26}(p, q) \leq d_e(p, q) \leq d_6(p, q) \leq 3 \cdot d_{26}(p, q)$ for all $p, q \in \mathbb{R}^3$, and $d_{26}(p, q) \leq d_{18}(p, q) \leq d_e(p, q)$ for all $p, q \in \mathbb{Z}^3$ such that $d_e(p, q) \neq \sqrt{3}$.

2D and 3D Geodesics

A sequence ρ of grid points $(p_0, p_1, \ldots, p_n)$ such that $p_0 = p, p_n = q$, and p_{i+1} is α-adjacent to p_i $(0 \leq i \leq n-1)$ is called an *α-path* of length n from p to q; p and q are called the *endpoints* of ρ.

An α-path is called an *α-geodesic* if no shorter α-path with the same endpoints exists.

Theorem 6 *The length of a shortest α-path from p to q is $d_\alpha(p,q)$.*

It follows that an α-path ρ of length n is an α-geodesic iff the d_α-distance between the endpoints of ρ is n.

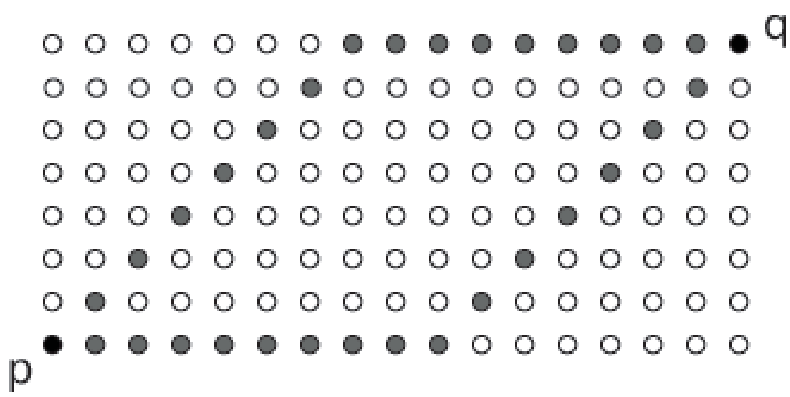

In Euclidean space, there is a unique shortest arc between any two points p and q, which is namely the straight line segment pq. In a grid, there can be many shortest α-paths between two grid points, and these paths need not be digital straight line segments.

2D and 3D Intrinsic Distances

If S is an α-connected set of grid points, for any $p, q \in S$, there exists an α-path $\rho = (p_0, p_1, \ldots, p_n)$ from $p_0 = p$ to $p_n = q$ such that the p_is are all in S.

The length $d_\alpha^S(p, q)$ of a shortest such path is called the *intrinsic α-distance* in S from p to q. The *α-eccentricity* $e_\alpha^S(p)$ is the maximum of all $d_\alpha^S(p, q)$, $q \in S$.

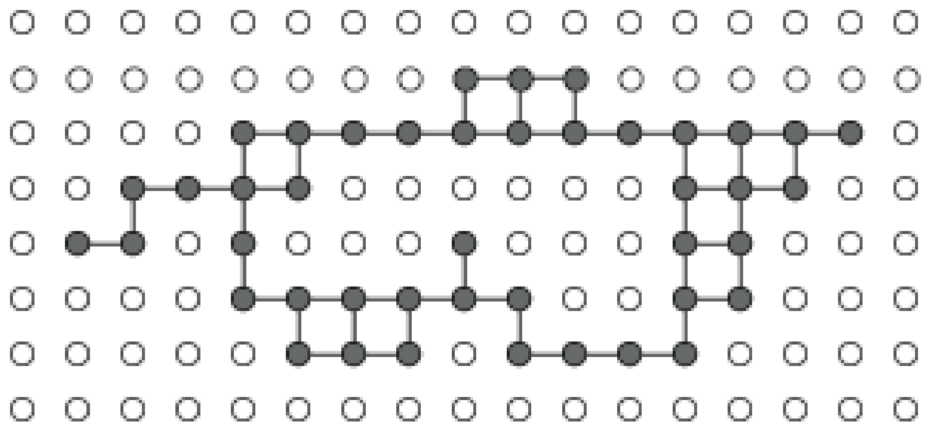

The *intrinsic α-diameter* $d_\alpha(S)$ of S is the maximum intrinsic α-distance between any of its points. This is equal to the maximum α-eccentricity of all $p \in S$.

If the intrinsic diameter of a connected set S of pixels or voxels is given by the intrinsic distance between $p, q \in S$, then both p and q are on the border of S.

The *intrinsic* α*-radius* $r_\alpha(S)$ of S is the minimum α-eccentricity of all $p \in S$.

p is an α*-central pixel or voxel* in S iff $e_\alpha^S(p) = r_\alpha(S)$.

The α*-center* $C_\alpha(S)$ of S is the set of all α-central pixels or voxels in S.

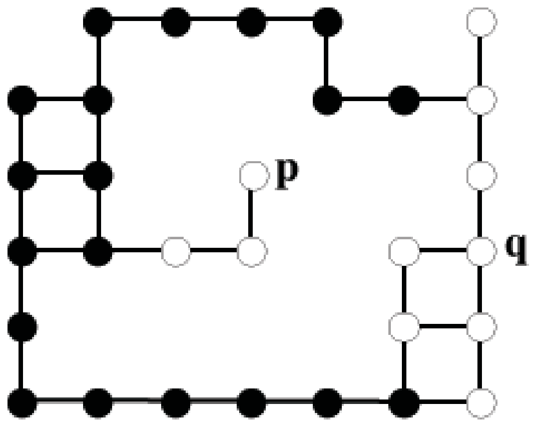

The intrinsic diameter of this 4-connected set is equal to the intrinsic 4-distance between pixels p and q. All the filled dots are 4-central pixels.

C. Jordan showed in 1869 that the centre of a tree is either a single node or a pair of adjacent nodes.

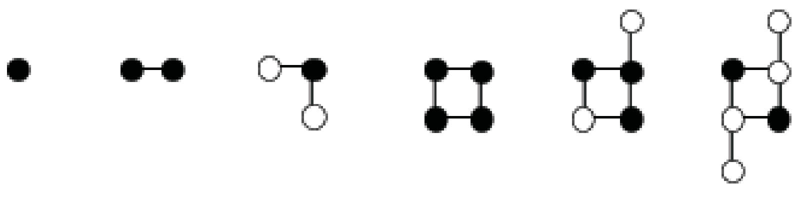

More examples of 4-centers. Centers of (larger) simply connected regions can be used to identify locations of these regions.

Coursework

Related material in textbook: Sections 1.2.1, 3.1.1, 3.1.2 (norms not needed), 3.1.4, 3.1.6, 3.2.1, and parts of 3.2.4. Do Exercise 2 on page 113.

A.7. [5 marks] Implement a program which calculates the intrinsic diameter (with respect to 4-adjacency) of any simply 4-connected set of pixels. (Hints: only a subset of border pixels needs to be considered, and in graph theory you learned, for example, about Dijkstra's algorithm).

Visualize results for 4-regions of varying shape complexity and size. Generate a diagram which summarizes the observed runtime of your algorithm depending on the size (in numbers of pixels) of used 4-regions.

Optional **[1 mark]**: you may extend your program such that it also allows to calculate the 4-center of a given simply 4-connected set of pixels.

Appendix: Center versus Centroid

For a finite set $S = \{p_1, \ldots, p_n\}$ of points in the plane, $p_i = (x_i, y_i)$ for $1 \leq i \leq n$, the *centroid* $p_c = (x_c, y_c)$ is defined by the mean of all coordinates:

$$x_c = \frac{1}{n}\sum_{i=1}^{n} x_i \quad \text{and} \quad y_c = \frac{1}{n}\sum_{i=1}^{n} y_i$$

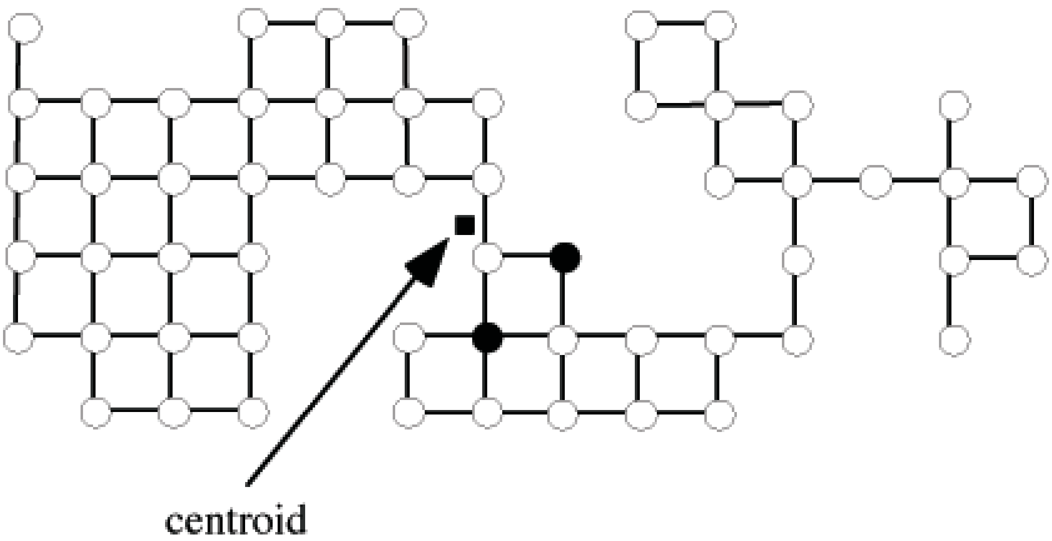

The centroid is unique, but in general not at a grid point position. The center may contain more than just one pixel, but all points in the center are at grid point position.