Computer Vision I

Bjoern Andres, Holger Heidrich

Machine Learning for Computer Vision
TU Dresden

Winter Term 2022/2023

Pixel classification

Digital image[1] $f\colon V \to C$

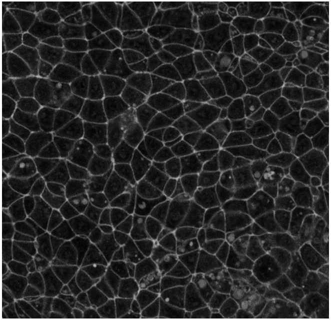

Classification $y\colon V \to \{0,1\}$

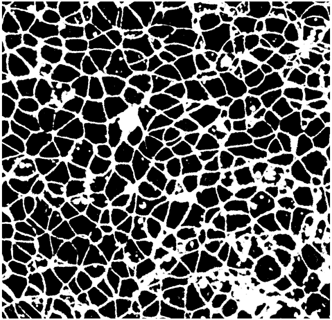

[1]By courtesy of Stephan Grill and his lab at the MPI of Molecular Cell Biology and Genetics.

Suppose we can construct a function $c\colon V \to \mathbb{R}$ wrt. a ditial image $f\colon V \to C$ in such a way that for any pixel $v \in V$:

- $c_v < 0$ if we consider $y_v = 1$ to be the right decision
- $c_v > 0$ if we consider $y_v = 0$ to be the right decision.

Definition 1. For any set V of pixels and any function $c\colon V \to \mathbb{R}$, the instance of the **trivial pixel classification problem** wrt. c has the form

$$\min_{y \in \{0,1\}^V} \sum_{v \in V} c_v \, y_v \tag{1}$$

In case the decision y_v for a pixel v depends on the color $f(v)$ of that pixel only, we can in principle

- construct a function $\xi\colon C \to \mathbb{R}$
- define $c_v = \xi(f(v))$ for any $v \in V$.

In practice, this task is supported by carefully designed GUIs.

In case the decision y_v for a pixel v depends on the colors of all pixels in a neighborhood $N(v) \subseteq V$ around v, we can in principle

- construct, for any pixel v, a function $\xi_v\colon C^{N(v)} \to \mathbb{R}$ that assigns a real number $\xi_v(f')$ to any coloring $f'\colon N(v) \to C$ of the neighborhood $N(v)$ of v
- define $c_v = \xi(f_{N(v)})$ for any $v \in V$.

In practice, this task is typically addressed by **machine learning**.

- In practice, solutions to the trivial pixel classification problem can be improved by exploiting **prior knowledge** about feasible combinations of decisions.
- Next, we consider prior knowledge saying that decisions at neighboring pixels $v, w \in V$ are more likely to be equal ($y_v = v_w$) than unequal ($y_v \neq y_w$).

Definition 2. For any pixel grid graph (V, E), any $c\colon V \to \mathbb{R}$ and any $c'\colon E \to \mathbb{R}_0^+$, the instance of the **smooth pixel classification problem** wrt. c and c' has the form

$$\min_{y \in \{0,1\}^V} \quad \underbrace{\sum_{v \in V} c_v \, y_v + \sum_{\{v,w\} \in E} c'_{\{v,w\}} \, |y_v - y_w|}_{\varphi(y)} \tag{2}$$

A naïve algorithm for this problem is local search with a transformation $T_v \colon \{0,1\}^V \to \{0,1\}^V$ that changes the decision for a single pixel, i.e., for any $y \colon V \to \{0,1\}$ and any $v, w \in V$:

$$T_v(y)(w) = \begin{cases} 1 - y_w & \text{if } w = v \\ y_w & \text{otherwise} \end{cases} \quad .$$

Initially, $y \colon V \to \{0,1\}$ and $W = V$
while $W \neq \emptyset$
    $W' := \emptyset$
    for each $v \in W$
        if $\varphi(T_v(y)) - \varphi(y) < 0$
            $y := T_v(y)$
            $W' := W' \cup \{w \in V \mid \{v,w\} \in E\}$
    $W := W'$

- So far, we have studied a local search algorithm for the smooth pixel classification problem.
- On the one hand, this algorithm is easy to implement and has straight-forward generalizations, e.g., to the case of more than two classes.
- On the other hand, it does not necessarily solve smooth pixel classification with two classes to optimality.
- Next, we will reduce the smooth pixel classification problem with two classes to the well-known **minimum st-cut problem** that can be solved exactly and efficiently.
- The notes are organized as follows
  - Definition of the minimum st-cut problem
  - Submodularity
  - Reduction of the smooth pixel classification problem

Definition 3. A 5-tuple $N = (V, E, s, t, \gamma)$ is called a **network** iff (V, E) is a directed graph and $s \in V$ and $t \in V$ and $s \neq t$ and $\gamma : E \to \mathbb{R}_0^+$.

The nodes s and t are called the **source** and the **sink** of N, respectively.

For any edge $e \in E$, γ_e is called the **capacity** of e in N.

Definition 4. Let (V, E) be a directed graph. Let $s \in V$ and $t \in V$ and $s \neq t$.

- $X \subseteq V$ is called an st-**cutset** of (V, E) iff $s \in X$ and $t \notin X$.
- $Y \subseteq E$ is called an st-**cut** of (V, E) iff there exists an st-cutset X such that $Y = \{vw \in E \mid v \in X \wedge w \notin X\}$.

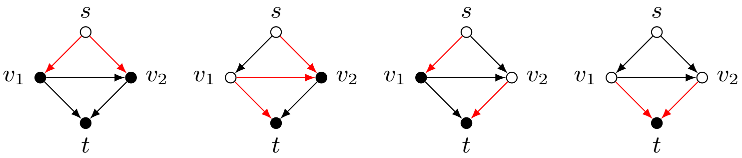

Definition 5. The instance of the **minimum** st-**cut problem** wrt. a network $N = (V, E, s, t, \gamma)$ is to

$$\begin{aligned} \min_{x \in \{0,1\}^V} \quad & \sum_{vw \in E} x_v \, (1 - x_w) \, \gamma_{vw} && (3) \\ \text{subject to} \quad & x_s = 1 && (4) \\ & x_t = 0 && (5) \end{aligned}$$

Definition 6. A **lattice** $(S, \preceq)$ is a set S, equipped with a partial order $\preceq$, such that any two elements of S have an infimum and a supremum wrt. $\preceq$.

Example. $(\{0,1\}^2, \preceq)$ with $\preceq := \{(s,t) \in S \times S \mid s_1 \leq t_1 \wedge s_2 \leq t_2\}$.

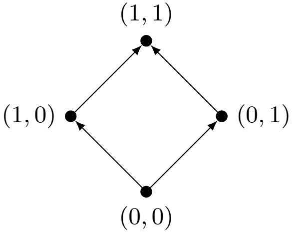

For any $s, t \in \{0,1\}^2$,

$$\sup(s,t) = (\max\{s_1, t_1\}, \max\{s_2, t_2\})$$
$$\inf(s,t) = (\min\{s_1, t_1\}, \min\{s_2, t_2\})$$

Definition 7. A function $f : S \to \mathbb{R}$ is called **submodular** wrt. a lattice $(S, \preceq)$ iff

$$\forall s, t \in S \qquad f(\inf(s,t)) + f(\sup(s,t)) \leq f(s) + f(t) \ . \tag{6}$$

Lemma 1. The sum of two submodular functions is submodular.

Lemma 2. For any $f : \{0,1\}^2 \to \mathbb{R}$, the following statements are equivalent.

1. f is is submodular wrt. the the lattice $(\{0,1\}^2, \preceq)$
2. $f(0,0) + f(1,1) \leq f(1,0) + f(0,1)$
3. The unique form

$$c_\emptyset + c_{\{1\}}x_1 + c_{\{2\}}x_2 + c_{\{1,2\}}x_1x_2$$

of f is such that $c_{\{1,2\}} \leq 0$.

Proof.

- $f(0,0) + f(1,1) \leq f(1,0) + f(0,1)$ is the only condition in

$$\forall s, t \in S \qquad f(\inf(s,t)) + f(\sup(s,t)) \leq f(s) + f(t)$$

  which is not generally true. Thus, (1.) is equivalent to (2.).

- We have

$$\begin{aligned} f(0,0) &= c_\emptyset \\ f(1,0) &= c_\emptyset + c_{\{1\}} \\ f(0,1) &= c_\emptyset \qquad\quad + c_{\{2\}} \\ f(1,1) &= c_\emptyset + c_{\{1\}} + c_{\{2\}} + c_{\{1,2\}} \quad . \end{aligned}$$

  Therefore,

$$c_{\{1,2\}} = f(1,1) - f(1,0) - f(0,1) + f(0,0)$$

  and thus, (2.) is equivalent to (3.).

Lemma 3. For every $f : \{0,1\}^2 \to \mathbb{R}$, there exist unique $a_0 \in \mathbb{R}$ and $a_1, a_{\bar{1}}, a_2, a_{\bar{2}}, a_{12}, a_{\bar{1}2} \in \mathbb{R}_0^+$ such that

$$a_1 a_{\bar{1}} = a_2 a_{\bar{2}} = a_{12} a_{\bar{1}2} = 0 \tag{7}$$

and

$$\begin{aligned} \forall x \in \{0,1\}^2 \quad f(x) = \quad & a_0 \\ & + a_1 x_1 + a_{\bar{1}}(1 - x_1) \\ & + a_2 x_2 + a_{\bar{2}}(1 - x_2) \\ & + a_{12} x_1 x_2 + a_{\bar{1}2}(1 - x_1) x_2 \quad . \end{aligned} \tag{8}$$

Proof.

- Comparison of (8) with the unique form

$$c_\emptyset + c_{\{1\}}x_1 + c_{\{2\}}x_2 + c_{\{1,2\}}x_1x_2$$

  yields

$$\begin{aligned} a_0 + a_{\bar{1}} + a_{\bar{2}} &= c_\emptyset \\ a_1 - a_{\bar{1}} &= c_{\{1\}} \\ a_2 - a_{\bar{2}} + a_{\bar{1}2} &= c_{\{2\}} \\ a_{12} - a_{\bar{1}2} &= c_{\{1,2\}} \end{aligned} \tag{9}$$

- By these equations (from bottom to top), (7) and c define a uniquely.

Lemma 4. For every **submodular** $f : \{0,1\}^2 \to \mathbb{R}$ and its unique coefficient $a_0 \in \mathbb{R}$ from Lemma 3,

$$\min_{x \in \{0,1\}^2} f_x - a_0 \tag{10}$$

is equal to the weight of a **minimum** st-**cut** in the graph below whose edge weights are the (unique, non-negative) coefficients from Lemma 3.

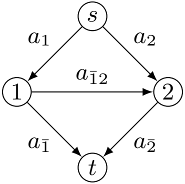

Moreover, f is minimal at $\hat{x} \in \{0,1\}^2$ iff $\{j \in \{1,2\} \mid \hat{x}_j = 0\}$ is a **minimum** st-**cutset** of the above graph.

Proof.

- Submodularity of f implies $a_{12} = 0$ in (9), by Lemma 2 and (7).
- Comparison of the four possible minima of f,

$$
\begin{aligned}
f(0,0) &= a_0 + a_{\bar{1}} + a_{\bar{2}} \\
f(1,0) &= a_0 + a_1 + a_{\bar{2}} \\
f(0,1) &= a_0 + a_{\bar{1}} + a_2 + a_{\bar{1}2} \\
f(1,1) &= a_0 + a_1 + a_2 + a_{12} \quad ,
\end{aligned}
$$

with the four possible minimum cuts below proves the Lemma.

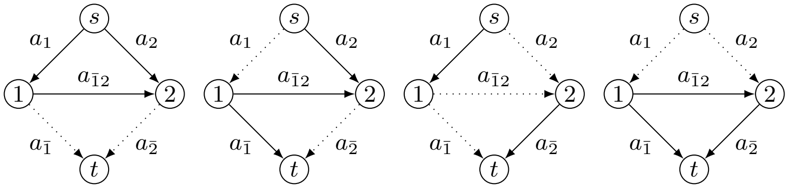

Definition 8. For any smooth pixel classification problem

$$\min_{y \in \{0,1\}^V} \quad \underbrace{\sum_{v \in V} c_v \, y_v + \sum_{\{v,w\} \in E} c'_{\{v,w\}} \, |y_v - y_w|}_{\varphi(y)} \tag{11}$$

the **induced minimum** st-**cut problem** is defined by the network (V', E', s, t, γ) such that $V' = V \cup \{s,t\}$,

$$\begin{aligned} E' =& \{(s,v) \in V'^2 \mid c_v > 0\} \cup \{(v,t) \in V'^2 \mid c_v < 0\} \\ & \cup \{(v,w) \in V'^2 \mid \{v,w\} \in E\} \end{aligned} \tag{12}$$

and $\gamma \colon E' \to \mathbb{R}_0^+$ such that

$$\forall (s,v) \in E' \colon \quad \gamma_{(s,v)} = c_v \tag{13}$$

$$\forall (v,t) \in E' \colon \quad \gamma_{(v,t)} = -c_v \tag{14}$$

$$\forall \{v,w\} \in E \colon \quad \gamma_{(v,w)} = \gamma_{(w,v)} = c'_{\{v,w\}} \quad . \tag{15}$$

Lemma 5. For any smooth pixel classification problem wrt. a pixel grid graph $G = (V, E)$ and the induced minimum st-cut problem with the network (V', E', s, t, γ), $\hat{y} : V \to \{0, 1\}$ is an optimal pixel classification iff $\{v \in V \mid \hat{y}_v = 0\}$ is an optimal st-cutset.

Proof (sketch). The function φ is submodular, by Lemma 1 and $c' > 0$.

The statement holds by Lemma 3 and the fact that for all $y \in \{0, 1\}^V$:

$$\varphi(y) = \sum_{v \in V} c_v \, y_v + \sum_{\{v,w\} \in E} c'_{\{v,w\}} \left(y_v(1 - y_w) + (1 - y_v) y_w \right) \quad .$$