

Combinatorial Continuous Maximum Flow*

Camille Couprie[†], Leo Grady[‡], Hugues Talbot[†], and Laurent Najman[†]

Abstract. Maximum flow (and minimum cut) algorithms have had a strong impact on computer vision. In particular, graph cut algorithms provide a mechanism for the discrete optimization of an energy functional which has been used in a variety of applications such as image segmentation, stereo, image stitching, and texture synthesis. Algorithms based on the classical formulation of max-flow defined on a graph are known to exhibit metrication artifacts in the solution. Therefore, a recent trend has been to instead employ a spatially *continuous* maximum flow (or the dual min-cut problem) in these same applications to produce solutions with no metrication errors. However, known fast continuous max-flow algorithms have no stopping criteria or have not been proved to converge. In this work, we revisit the continuous max-flow problem and show that the analogous discrete formulation is *different* from the classical max-flow problem. We then apply an appropriate combinatorial optimization technique to this combinatorial continuous max-flow (CCMF) problem to find a null-divergence solution that exhibits no metrication artifacts and may be solved exactly by a fast, efficient algorithm with provable convergence. Finally, by exhibiting the dual problem of our CCMF formulation, we clarify the fact, already proved by Nozawa in the continuous setting, that the max-flow and the total variation problems are not always equivalent.

Key words. total variation minimization, interior point methods, discrete calculus, image segmentation, minimal surfaces

AMS subject classifications. 35J05, 35L05, 35Q90, 35Q93, 65K10, 65K15, 68U10, 90C25, 90C35, 90C51

DOI. 10.1137/100799186

1. Introduction. Optimization methods have been used to address a wide variety of problems in computer vision. The early optimization approaches were formulated in terms of active contours and surfaces [29] and then later level sets [41]. These formulations were used to optimize energies of varied sophistication (e.g., using regional, texture, motion, or contour terms [36]) but generally converged to a local minimum, generating results that were sensitive to initial conditions and noise levels. Consequently, more recent focus has been on energy formulations (and optimization algorithms) for which a global optimum can be found.

Such energy formulations typically include a term which minimizes the boundary length (or surface area in three dimensions) of a region or the total variation of a scalar field in addition to a data term and/or hard constraints. In this paper, we focus on image segmentation as the example application on which to base our exposition. Indeed, segmentation has played a prominent (and early) role in the development of the use of global optimization methods in computer vision and often forms the basis of many other applications [7, 28, 14].

*Received by the editors June 16, 2010; accepted for publication (in revised form) May 31, 2011; published electronically September 15, 2011.

<http://www.siam.org/journals/siims/4-3/79918.html>

[†]Laboratoire d'Informatique Gaspard-Monge, Université Paris-Est, Equipe A3SI, ESIEE Paris, 93160 Noisy-le-Grand, France (c.couprie@esiee.fr, h.talbot@esiee.fr, l.najman@esiee.fr).

[‡]Department of Imaging and Visualization, Siemens Corporate Research, Princeton, NJ 08540 (leo.grady@siemens.com).

Graph-based approaches. The max-flow/min-cut problem on a graph is a classical problem in graph theory, for which the earliest solution algorithm goes back to Ford and Fulkerson [17]. Initial methods for global optimization of the boundary length of a region formulated the energy on a graph and relied on max-flow/min-cut methods for solution [5, 30]. It was soon realized that these methods introduced a metrication error. Metrication errors are clearly visible when gradient information is weak or missing, for example in the case of medical image segmentation or in some materials science applications like electron tomography, where weak edges are unavoidable features of the imaging modality. Metrication errors are also far more obvious in three dimensions than in two dimensions and increasingly so, as the dimensionality increases (see, for example, the three-dimensional (3D) lung example in [1]). Furthermore, metrication artifacts are even more present in other applications such as surface reconstruction and rendering, where the artifacts are a lot worse than in image segmentation. Various solutions for metrication errors were proposed. One solution involved the use of a highly connected lattice with a specialty edge weighting [6], but the large number of graph edges required to implement this solution could cause memory concerns when implemented for large two-dimensional (2D) or 3D images [1].

Continuous approaches. To avoid the metrication artifacts without increasing memory usage, one trend in recent years has been to pursue spatially *continuous* formulations of the boundary length and the related problem of total variation [12, 1, 13]. Historically, a continuous max-flow (and dual min-cut problem) was formulated by Iri [27] and then Strang [43]. Strang's continuous formulation provided an example of a *continuization* (as opposed to *discretization*) of a classically discrete problem, but it was not associated with any algorithm. Work by Appleton and Talbot [1] provided the first PDE-based algorithm to find Strang's continuous max-flows and therefore optimal min-cuts. This same algorithm was also derived by Unger et al. [46] from the standpoint of minimizing continuous total variation. Adapted to image segmentation, this algorithm is shown to be equivalent [45] to the Appleton–Talbot algorithm and has been demonstrated to be fast when implemented on massively parallel architectures. Unfortunately, this algorithm has no stopping criteria and has not been proved to converge. Works by Pock et al. [38], Zach et al. [48], and Chambolle, Cremers, and Pock [10] present different algorithms for optimizing comparable energies for solving multilabel problems, but again, those algorithms are not proved to converge. Some other works have been presenting provably converging algorithms, but with relatively slow convergence speed. For example, Pock and Chambolle introduce in [11] a general saddle point algorithm that may be used in various applications, but it needs half an hour to segment a 350×450 image on a CPU and still more than a dozen seconds on a GPU (details are given in the experiments section).

Links and differences with total variation minimization. Strang [44] has shown the continuous max-flow problem for the l_2 norm to be the dual of the total variation (TV) minimization problem under some assumptions. The problem of total variations, introduced successively in computer vision by Shulman and Hervé [42] and later Rudin, Osher, and Fatemi [40] as a regularizing criterion for optical flow computation and then image denoising, has been shown to be quite efficient for smoothing images without affecting contours. Moreover, a major advantage of TV is that it is a convex problem. Thus, a straightforward algorithm such as gradient descent may be applied to find a minimum solution. However, there is a need for faster methods, and significant progress has been achieved recently with primal-dual

approaches [12], Nesterov's algorithm [33], and split Bregman/Douglas–Rachford methods [21, 49]. Most methods minimizing TV focus on image filtering as an application, and even if those methods are remarkably fast in denoising applications, segmentation problems require a lot more iterations for those algorithms to converge. As stated previously, continuous max-flow is dual with TV *in the continuous setting under restrictive regularity conditions* [34]. In fact, Nozawa [35] showed that there is a duality gap between weighted TV and weighted max-flow under some conditions in the continuous domain. Thus, it is important to note that weighted TV problems are not equivalent to the weighted maximum flow problem studied in this paper. Furthermore, many works assume that the continuous-domain duality holds algorithmically, but we show later in this paper that at least in the combinatorial case this is not true.

The previous works for solving the max-flow problem illustrate two difficulties with continuous-based formulation, which are (1) the discretization step, which is necessary for deriving algorithms, but may break continuous-domain properties; and (2) the convergence, both of the underlying continuous formulation and the associated algorithm, which itself depends on the discretization. It is very well known that even moderately complex systems of PDEs may not converge, and even existence of solutions is sometimes not obvious [16]. Even when existence and convergence proofs both exist, sometimes algorithmic convergence may be slow in practice. For these reasons, combinatorial approaches to the maximal flow and related problems are beginning to emerge.

Combinatorial approaches. Discrete calculus [26, 25] has been used in recent years to produce a combinatorial reformulation of continuous problems onto a graph in such a manner that the solution behaves analogously to the continuous formulation (see, e.g., [15, 22]).

In particular, Lippert presented in [32] a combinatorial formulation of an isotropic continuous max-flow problem on a planar lattice, making it possible to obtain a provably optimal solution. However, in Lippert's work, parameterization of the capacity constraint is tightly coupled to the 4-connected squared grid, and the generalization to higher dimensions seems nontrivial. Furthermore it involves the multiplication of the number of capacity constraints by the degree of each node, thus increasing the dimension of the problem. This formulation did not lead to a fast algorithm. The author compared several general solvers, quoting an hour as their best time for computing a solution on a 300×300 lattice.

Motivation and contributions. In this paper, we pursue a combinatorial reformulation of the continuous max-flow problem initially formulated by Strang, which we term *combinatorial continuous maximum flow* (CCMF). Viewing our contribution differently, we adopt a discretization of continuous max-flow as the primary problem of interest, and then we apply fast combinatorial optimization techniques to solve the discretized version.

Our reformulation of the continuous max-flow problem produces a (divergence-free) flow problem defined on an arbitrary graph. Strikingly, CCMF are not equivalent to the discretization of continuous max-flows produced by Appleton and Talbot or that of Lippert (if for no other reason than the fact that CCMF is defined on an arbitrary graph). Moreover, *CCMF is not equivalent to the classical max-flow on a graph*. In particular, we will see that the difference lies in the fact that capacity constraints in classical max-flow restrict flows along graph edges while the CCMF capacity constraints restrict the total flow passing through the nodes.

The CCMF problem is convex. We deduce an expression of the dual problem, which allows us to employ a primal-dual interior point method to solve it; e.g., interior point methods have

been used for graph cuts [2] and second order cone problems in general [20]. The CCMF problem has several desirable properties, including the following:

1. The solution to CCMF on a 4-connected lattice avoids the metrication errors. Therefore, the gridding error problems may be solved without the additional memory required to process classical max-flow on a highly-connected graph.
2. In contrast to continuous max-flow algorithms of Appleton and Talbot (AT-CMF) and equivalent, the solution to the CCMF problem can be computed with guaranteed convergence in practical time.
3. The CCMF problem is formulated on an arbitrary graph, which enables the algorithm to incorporate nonlocal edges [8, 15, 24] or to apply it to arbitrary clustering problems defined on a graph.
4. The algorithm for solving the CCMF problem is fast, easy to implement, compatible with multiresolution grids, and straightforward to parallelize.

Our computation of the CCMF dual further reveals that duality between TV minimization and maximum flow does not hold for CCMF and combinatorial total variation (CTV). The comparison between those two combinatorial problems is motivated by several interests:

1. For clarification: In the continuous domain, the duality between TV and MF holds under some regularity constraints. In the discrete anisotropic domain (i.e. graph cuts), this duality always holds. However, in the isotropic weighted discrete case, i.e., in the CTV, AT-CMF, or CCMF case, we are not aware of any discretization such that the duality holds. It is not obvious to realize that the duality does not hold in the discrete setting even though it may in the continuous case. We present clearly the differences between the two problems, theoretically and in term of results.
2. To expose links between CCMF and CTV: We prove that the weak duality holds between the two problems.
3. For efficiency: In image segmentation, the fastest known algorithms to optimize CTV are significantly slower than CCMF. Therefore, there is reason to believe that CCMF may be used to efficiently optimize energies for which TV has been previously shown to be useful (e.g., image denoising).

In the next section, we review the formulation of continuous max-flow, derive the CCMF formulation and its dual, and then provide details of the fast and provably convergent algorithm used to solve the new CCMF problem.

2. Method. Our combinatorial reformulation of the continuous maximum flow problem leads to a formulation on a graph which is *different* from the classical max-flow algorithm. Before proceeding to our formulation, we review the continuous max-flow and the previous usage of this formulation in computer vision.

2.1. The continuous max-flow (CMF) problem. First introduced by Iri [27], Strang presents in [43] an expression of the continuous maximum flow problem

$$(2.1) \quad \begin{aligned} & \max F_{st} \\ \text{s.t. } & \nabla \cdot \vec{F} = 0, \\ & \|\vec{F}\| \leq g. \end{aligned}$$

Here F_{st} is the total amount of flow going from the source location s to the sink location t , \vec{F} is the flow, and g is a scalar field representing the local metric distortion. The solution to this problem is the exact solution of the geodesic active contour (GAC) (or surface) formulation [9]. In order to solve the problem (2.1), the Appleton–Talbot algorithm (AT-CMF) [1] solves the following PDE system:

$$(2.2) \quad \begin{aligned} \frac{\partial P}{\partial \tau} &= -\nabla \cdot \vec{F}, \\ \frac{\partial \vec{F}}{\partial \tau} &= -\nabla P \\ \text{s.t. } \|\vec{F}\| &\leq g. \end{aligned}$$

Here P is a potential field similar to the excess value in the Push-Relabel maximum flow algorithm [19]. AT-CMF is effectively a simple continuous computational fluid dynamics (CFD) simulation with nonlinear constraints. It uses a forward finite-difference discretization of the above PDE system subject to a Courant–Friedrichs–Lewy (CFL) condition, also seen in early level-set methods. At convergence, the potential function P approximates an indicator function, with 0 values for the background labels and 1 for the foreground, and the flow \vec{F} has zero-divergence almost everywhere. However, there is no guarantee of convergence for this algorithm and, in practice, many thousands of iterations can be necessary to achieve a binary P , which can be very slow.

Although Appleton–Talbot is a continuous approach, applying this algorithm to image processing involves a discretization step. The capacity constraint $\|\vec{F}\| \leq g$ is interpreted as $\max F_{x(i)}^2 + \max F_{y(i)}^2 \leq g_i^2$, with $F_{x(i)}$ the outgoing flow of edges linked to node i along the x axis, and $F_{y(i)}$ the outgoing flow linked to node i along the y axis. We notice that the weights are associated with point locations (pixels), which will correspond later to a node-weighted graph.

In the next section, we use the operators of discrete calculus to reformulate the CMF problem on a graph and show the surprising result that the graph formulation of the continuous max-flow leads to a problem different from the classical max-flow problem on a graph.

2.2. A discrete calculus formulation. Before establishing the discrete calculus formulation of the CMF problem, we specify our notation. A graph consists of a pair $G = (V, E)$ with vertices $v \in V$ and edges $e \in E \subseteq V \times V$. A transport graph $G(V, E)$ comprises two additional nodes, named a source s and a sink t , and additional edges linking some nodes to the sink and some to the source. Including the source and the sink nodes, the cardinalities of G are given by $n = |V|$ and $m = |E|$. An edge, e , spanning two vertices, v_i and v_j , is denoted by e_{ij} . In this paper we deal with weighted graphs that include weights on both the edges and the nodes. An edge weight is a value assigned to each edge which may be viewed as a capacity in the context of a maximum flow problem. The weight of an edge, e_{ij} , is denoted by \tilde{g}_{ij} . In this work, we assume that $\tilde{g}_{ij} \in \mathbb{R}$ and $\tilde{g}_{ij} > 0$, and use \tilde{g} to denote the vector of \mathbb{R}^m containing the \tilde{g}_{ij} for every edge e_{ij}

of G . In addition to edge weights, we may also assign weights to nodes. The weight of node v_i is denoted by g_i . In this work, we also assume that $g_i \in \mathbb{R}$ and $g_i > 0$. We use g to denote the vector of \mathbb{R}^n containing the g_{ij} for every edge e_{ij} of G . We define a flow through edge e_{ij} as F_{ij} , where $F_{ij} \in \mathbb{R}$, and use the vector $F \in \mathbb{R}^m$ to denote the flows through all edges in the graph. Each edge is assumed to be oriented, such that a positive flow on edge e_{ij} indicates the direction of flow from v_i to v_j , while a negative flow indicates the direction of flow from v_j to v_i .

The incidence matrix is a key operator for defining a combinatorial formulation of the CMF problem. Specifically, the incidence matrix $A \in \mathbb{R}^{m \times n}$ is known to define the discrete calculus analogue of the gradient, while A^T is known to define the discrete calculus analogue of the divergence (see [25] and the references therein). The incidence matrix maps functions on nodes (a scalar field) to functions on edges (a vector field) and may be defined as

$$(2.3) \quad A_{e_{ij}v_k} = \begin{cases} +1 & \text{if } i = k, \\ -1 & \text{if } j = k, \\ 0 & \text{otherwise} \end{cases}$$

for every vertex v_k and edge e_{ij} . In our formulation of continuous max-flow, we use the expression $|A|$ to denote the matrix formed by taking the absolute value of each entry individually.

Given these definitions, we now produce a discrete (combinatorial) version of the continuous max-flow of (2.1) on a transport graph. As in [25, 15, 22], the continuous vector field indicating flows may be represented by a vector on the edge set, F . Additionally, the combinatorial divergence operator allows us to write the first constraint in (2.1) as $A^T F = 0$. The second constraint in (2.1) involves the comparison of the pointwise norm of a vector field with a scalar field. Therefore, we can follow [25, 15, 22] to define the pointwise ℓ_2 norm of the flow field F as $\sqrt{|A^T|F^2}$. In our notation here, as in the rest of the paper, $F^2 = F \cdot F$ denotes an elementwise product, “ \cdot ” denoting the Hadamard (elementwise) product between the vectors, and the square root of a vector is also here and in the rest of the paper the vector composed of the square roots of every element.

Putting these pieces together, we obtain

$$(2.4) \quad \begin{aligned} & \max F_{st} \\ \text{s.t. } & A^T F = 0, \\ & |A^T|F^2 \leq g^2. \end{aligned}$$

Compare this formulation to the classical max-flow problem on a graph, given in our notation as

$$(2.5) \quad \begin{aligned} & \max F_{st} \\ \text{s.t. } & A^T F = 0, \\ & |F| \leq \tilde{g}. \end{aligned}$$

By comparing these formulations of the traditional max-flow with our combinatorial formulation of the continuous max-flow, it is apparent that the key difference between the classical formulation and our combinatorial continuous max-flow is in the capacity constraints.




Figure 1. The difference between classical max-flow on a graph with CCMF on a graph is that classical max-flow uses edge-weighted capacities while CCMF uses node-weighted capacities. This difference is manifest in the different solutions obtained for both algorithms and the algorithms required to find a solution. Specifically, the solution to the CCMF problem on a lattice does not exhibit metrication bias.

In both formulations, the flow is defined through edges, but in the classical case the capacity constraint restricts flow through an edge while the CCMF formulation restricts the amount of flow passing through a node by taking in account the neighboring flow values. This contrast-weighting applied to nodes (pixels) has been a feature of several algorithms with a continuous formulation, including GACs [9], CMF [1], TVSeg [46], and graph cuts [11]. On the other hand, the problem defined in (2.4) is defined on an arbitrary graph in which contrast weights (capacities) are almost always defined on edges. The node-weighted capacities fit Strang’s formulation of a scalar field of constraints in the continuous max-flow formulation, and therefore our graph formulation of Strang’s formulation carries over these same capacities. The most important point here is not having expressed the capacity on the nodes of the graph but rather how the flow is enforced to be bounded by the metric in each node. Bounding the norm of neighboring flow values in each node simulates a closer behavior of the continuous setting than if no contribution of neighboring flow were made as in the standard max-flow problem. Figure 1 illustrates the relationship of the edge capacity constraints in the classical max-flow problem to the node capacity constraints in the CCMF formulation. The null-divergence constraint is also essential in our formulation since it permits us to obtain constant partitions almost everywhere in the dual solution introduced in the next section, in particular binary ones.

In the context of image segmentation, the vector g varies inversely to the image gradient. We propose using, as in [46],

$$(2.6) \quad g = \exp(-\beta \|\nabla I\|_2),$$

where I indicates the image intensity. For simplicity, this weighting function is defined for greyscale images, but g could be used to penalize changes in other relevant image quantities, such as color or texture. Before addressing the solution of the CCMF problem, we consider the dual of the CCMF problem and, in particular, the sense in which it represents a minimum cut. Since the cut weights are present on the nodes rather than the edges, we must expect the minimum cut formulation to be different from the classical minimum cut on a graph.

2.3. The CCMF dual problem. Classically, the max-flow problem is dual to the min-cut problem, allowing a natural geometric interpretation of the objective function. In order to provide the same interpretation, we now consider the dual problem to the CCMF and show that we optimize a node-weighted minimum cut. In the following proposition, we denote the element-wise quotient of two vectors $u = [u_1, \dots, u_k]$ and $v = [v_1, \dots, v_k]$ by $u \cdot / v = [u_1/v_1, \dots, u_k/v_k]$. We also denote by $\mathbf{1}^n$ a unit vector $[1, \dots, 1]$ of size n and recall that the square exponent v^2 of a vector v represents the resulting vector of the elementwise multiplication $v \cdot v$.

Proposition 2.1. *In a transport graph G with m edges and n nodes, we define a vector c of \mathbb{R}^m , composed of zeros except for the element corresponding to the source/sink edge, which is 1. Let λ and ν be two vectors of \mathbb{R}^n . The CCMF problem*

$$(2.7) \quad \begin{aligned} \max_{F \in \mathbb{R}^m} \quad & c^T F \\ \text{s.t.} \quad & A^T F = 0, \\ & |A^T| F^2 \leq g^2 \end{aligned}$$

has for a dual

$$(2.8) \quad \begin{aligned} \min_{\lambda \in \mathbb{R}^n, \nu \in \mathbb{R}^n} \quad & \lambda^T g^2 + \frac{1}{4} \left(\mathbf{1}^n \cdot / (|A|\lambda) \right)^T \left((c + A\nu)^2 \right) \\ \text{s.t.} \quad & \lambda \geq 0, \end{aligned}$$

equivalently written in (2.15), and the optimal solution (F^*, λ^*, ν^*) verifies

$$(2.9) \quad \max_F c^T F = c^T F^* = 2\lambda^{*T} g^2$$

and the n following equalities:

$$(2.10) \quad \lambda^* \cdot |A^T| \left((c + A\nu^*) \cdot / (|A|\lambda^*) \right)^2 = 4\lambda^* \cdot g^2.$$

Proof. The Lagrangian of (2.7) is

$$L(F, \lambda, \nu) = (c^T + \nu^T A^T) F + \lambda^T |A^T| F^2 - \lambda^T g^2,$$

with $\lambda \in \mathbb{R}^n$ and $\nu \in \mathbb{R}^n$ two Lagrange multipliers. We have to find F in the dual function such that $\nabla_F L(F, \lambda, \nu) = 0$:

$$(2.11) \quad \nabla_F L(F, \lambda, \nu) = c + A\nu + 2|A|\lambda \cdot F = 0 \Leftrightarrow F = (c + A\nu) \cdot / (-2|A|\lambda).$$

Substituting F in the Lagrangian function, we obtain

$$L(\lambda, \nu) = (c^T + \nu^T A^T) \left((c + A\nu) \cdot / (-2|A|\lambda) \right) + \lambda^T |A^T| \left((c + A\nu) \cdot / (-2|A|\lambda) \right)^2 - \lambda^T g^2.$$

The Lagrangian may be simplified as

$$L(\lambda, \nu) = - \left(\mathbf{1}^n \cdot / (4|A|\lambda) \right)^T \left((c + A\nu)^2 \right) - \lambda^T g^2.$$

The dual of (2.7) is also

$$(2.12) \quad \begin{aligned} \min_{\lambda, \nu} \quad & \lambda^T g^2 + \frac{1}{4} \left(\mathbf{1}^n \cdot /(|A|\lambda) \right)^T \left((c + A\nu)^2 \right) \\ \text{s.t.} \quad & \lambda \geq 0. \end{aligned}$$

Furthermore, the KKT optimality conditions give

$$(2.13) \quad \lambda^* \cdot \left(|A^T| F^{*2} - g^2 \right) = 0.$$

Using the expression of F^* from (2.11), and substituting it in (2.13), we obtain

$$\lambda^* \cdot |A^T| \left((c + A\nu^*) \cdot /(|A|\lambda^*) \right)^2 = 4\lambda^* \cdot g^2.$$

Taking the sum of the right-hand side and the left-hand side,

$$(2.14) \quad (c + A\nu^*)^2 \cdot /(|A|\lambda^*) = 4\lambda^{*T} g^2.$$

Finally, if we substitute (2.14) in the dual expression (2.8), we have

$$\max_F c^T F = c^T F^* = 2\lambda^{*T} g^2. \quad \blacksquare$$

The expression of the CCMF dual may be written in summation form as

$$(2.15) \quad \begin{aligned} \min_{\lambda, \nu} \sum_{v_i \in V} \quad & \overbrace{\lambda_i g_i^2}^{\text{weighted cut}} + \overbrace{\frac{1}{4} \sum_{e_{ij} \in E \setminus \{s,t\}} \frac{(\nu_i - \nu_j)^2}{\lambda_i + \lambda_j}}^{\text{smoothness term}} + \overbrace{\frac{1}{4} \frac{(\nu_s - \nu_t - 1)^2}{\lambda_s + \lambda_t}}^{\text{source/sink enforcement}} \\ \text{s.t.} \quad & \lambda_i \geq 0 \quad \forall i \in V. \end{aligned}$$

Interpretation. The optimal value λ^* is a weighted indicator of the saturated vertices (a vertex v_i is saturated if $|A^T|_i F^2 = g_i^2$, where $|A^T|_i$ indicates the i th row of $|A^T|$):

$$(2.16) \quad \lambda^*(v_i) \begin{cases} > 0 & \text{if } |A^T|_i F^2 = g(v_i)^2, \\ = 0 & \text{otherwise.} \end{cases}$$

The variables ν_s and ν_t are not constrained to be set to 0 and 1; only their difference is constrained to be equal to one. Thus their values range between a constant and a constant plus one. Let us call this constant δ , and without loss of generality one can consider that $\delta = 0$. The term ν is at optimality a weighted indicator of the source/sink/saturated vertices partition:

$$\nu^*(v_i) = \begin{cases} 0 + \delta & \text{if } v_i \in S, \\ \text{a number between } (0 + \delta) \text{ and } (1 + \delta) & \text{if } |A^T|_i F^2 = g(v_i)^2, \\ 1 + \delta & \text{if } v_i \in T. \end{cases}$$

The expression (2.9) of the CCMF dual shows that the problem is equivalent to finding a minimum weighted cut defined on the nodes.

Finally, the “weighted cut” is recovered in (2.15), and the “smoothness term” is compatible with large variations of ν at the boundary of objects because of a large denominator (λ) in the contour area. An illustration of optimal λ and ν on an image is shown in Figure 2.




Figure 2. The dual problem to CCMF is a node-weighted minimum cut in which the variable λ is a weighted indicator vector labeling boundary nodes and the variable ν is a nearly binary vector indicating the source/sink regions. As a result, the contours of ν are slightly blurry. This is due to the equilibrium effect between the two dual variables. In practice, as λ is nonzero only in presence of a contour, ν is binary almost everywhere, except on a very thin line.

2.4. Solving the CCMF problem. When considering how to optimize the CCMF problem (2.4), the first key observation is that the constraints bound a convex feasible region.

The real valued functions $f_i : \mathbb{R}^m \rightarrow \mathbb{R}$ defined for every node $i = 1, \dots, n$ as $f_i(F) = |A^T|_i F^2 - g_i^2$ are nonnegative quadratic, so they are convex functions. We define the vector $f(F)$ of \mathbb{R}^n as $f(F) = [f_1(F), \dots, f_n(F)]^T$.

Since the constraints are convex, the CCMF problem may be solved with a fast primal-dual interior point method (see [4]), which we now review in the specific context of CCMF.

The primal-dual interior point (PDIP) algorithm iteratively computes the primal F and dual λ, ν variables so that the Karush–Kuhn–Tucker (KKT) optimality conditions are satisfied. This algorithm solves the CCMF problem (2.7) by applying Newton’s method to a sequence of a slightly modified version of the KKT conditions. The steps of the PDIP procedure are given in Algorithm 1. Specifically for CCMF, the system $M\Delta_y = r$ system may be written as

$$(2.17) \quad \begin{bmatrix} \sum_{i=1}^n \lambda_i \nabla^2 f_i(F) & Df(F)^T & A \\ -\text{diag}(\lambda) Df(F) & -\text{diag}(f(F)) & 0 \\ A^T & 0 & 0 \end{bmatrix} \begin{bmatrix} \Delta_F \\ \Delta_\lambda \\ \Delta_\nu \end{bmatrix} = - \begin{bmatrix} r_d = c + Df(F)^T \lambda + A\nu \\ r_c = -\text{diag}(\lambda) f(F) - (1/t) \\ r_p = A^T F \end{bmatrix}$$

with r_d , r_c , and r_p representing the dual, central, and primal residuals. Additionally, the derivatives are given by

$$\nabla f_i(F) = 2|A^T|_i \cdot F, \quad Df(F) = \begin{bmatrix} \nabla f_1(F)^T \\ \vdots \\ \nabla f_n(F)^T \end{bmatrix}, \quad \nabla^2 f_i(F) = \text{diag} \begin{bmatrix} 2|A^T|_{i1} \\ \vdots \\ 2|A^T|_{im} \end{bmatrix},$$

with “ \cdot ” denoting the Hadamard (elementwise) product between the vectors.

ALGORITHM 1. Primal dual interior point algorithm

Data: $F \in \mathbb{R}^m = 0, \lambda \in (\mathbb{R}_+^*)^n, \nu \in \mathbb{R}^n, \mu > 1, \epsilon > 0.$

Result: F solution to the CCMF problem (2.7) such as F_{st} is maximized under the divergence-free and capacity constraints, and ν, λ solution to the CCMF dual problem (2.9).

repeat

- 1. Compute the surrogate duality gap $\hat{\eta} = -f(F)^T \lambda$ and set $t = \mu n / \hat{\eta}.$
- 2. Compute the primal-dual search direction Δ_y such as $M \Delta_y = r.$
- 3. Determine a step length $s > 0$ and set $y = y + s \Delta_y$ ($y = [F, \lambda, \nu]^T$).

until $\|r_p\|_2 \leq \epsilon, \|r_d\|_2 \leq \epsilon,$ and $\hat{\eta} \leq \epsilon.$

Consequently, the primary computational burden in the CCMF algorithm is the linear system resolution required by (2.17). In the result section we present execution times obtained in our experiments using a conventional CPU implementation.

Observe that, although this linear system is large, it is very sparse and is straightforward to parallelize on a GPU [3, 31, 23], for instance, using an iterative GPU solver. If it does not necessarily imply a faster solution, the asymptotic complexity of modern iterative and modern direct solvers is about the same, and, in our experience, there has been a strong improvement in the performance of a linear system solve when going from a direct solver CPU solution to a conjugate gradient solver GPU solution, especially as the systems get larger.

3. Comparison between CCMF and existing related approaches.

3.1. Min-cut. As detailed in section 2.3, there is a difference between the CCMF and the classic max-flow formulation in the capacity constraint. Therefore, the dual problems are different. As proved by Ford and Fulkerson [17], and formalized, for example, in [2], the min-cut problem writes with our notation of section 2.3 as

$$(3.1) \quad \min_u \tilde{g}^T |Au + c|.$$

We note that the ℓ_1 norm of the gradient of the solution u of the above min-cut expression turns into an ℓ_2 norm of the gradient of the solution ν in the CCMF dual (2.8).

3.2. Primal-dual total variations. Unger, Pock, and Bischof [45] proposed optimizing the following primal-dual TV formulation:

$$(3.2) \quad \min_u \max_F \sum_{i,j} (u_{i+1,j} - u_{i,j}) F_{i,j}^1 + (u_{i,j+1} - u_{i,j}) F_{i,j}^2 + \text{data fidelity}$$

$$\text{s.t. } \sqrt{(F_{i,j}^1)^2 + (F_{i,j}^2)^2} \leq g_{i,j}.$$

We can note that the CCMF flow capacity constraint in a 4-connected lattice is similar to the constraint in (3.2) with a slight modification and would be with finite-element notation

$$(3.3) \quad \sqrt{(F_{i-1,j}^1)^2 + (F_{i,j-1}^2)^2 + (F_{i,j}^1)^2 + (F_{i,j}^2)^2} \leq g_{i,j}.$$

In contrast to this finite-element discretization, the provided CCMF formulation of continuous max-flow is defined on arbitrary graphs. Furthermore, we note that the optimization procedure employed to solve (3.2) generalizes the Appleton–Talbot algorithm.

3.3. Combinatorial total variations. We now compare the dual CCMF problem with the combinatorial total variation (CTV) problem. We note that the CCMF dual is different from CTV and discuss the weak duality of the two problems.

We recall the continuous expression of TV given by Strang:

$$(3.4) \quad \begin{aligned} \min_u \quad & \iint_{\Omega} g \|\nabla u\| \, dx \, dy \\ \text{s.t.} \quad & \int_{\Gamma} u f \, ds = 1, \end{aligned}$$

where Ω represents the image domain and Γ the boundary of Ω . The source and sink membership is represented by f , such that $f(x, y) > 0$ if (x, y) belongs to the sink, $f(x, y) < 0$ if (x, y) belongs to the source, and $f = 0$ otherwise. Considering a transport graph G and a vector u defined on the nodes, this continuous problem may be written with combinatorial operators as

$$(3.5) \quad \begin{aligned} \min_{u \in \mathbb{R}^n} \quad & g^T \sqrt{|A^T|(Au)^2} \\ \text{s.t.} \quad & u_s - u_t = 1, \end{aligned}$$

where the square root operator of a vector $v = [v_1, \dots, v_k]$ is an elementwise square root operator $\sqrt{v} = [\sqrt{v_1}, \dots, \sqrt{v_k}]$. Another way to write the same problem is

$$(3.6) \quad \begin{aligned} \min_u \quad & \sum_{v_i \in V} g_i \sqrt{\sum_{e_{ij} \in E} (u_i - u_j)^2} \\ \text{s.t.} \quad & u_s - u_t = 1. \end{aligned}$$

We note that in these equations the capacity g must be defined on the vertices. Although we describe this energy as “combinatorial total variation,” due to its derivation in discrete calculus terms, it is important to note that this formulation is very similar to the discretization which has appeared previously in the literature from Gilboa and Osher [18] and Chambolle [12] (if we allow $g = 1$ everywhere).

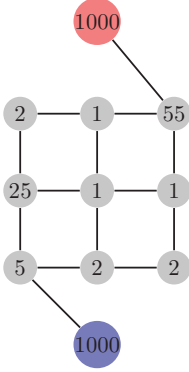
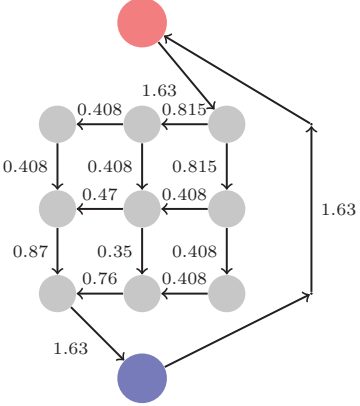
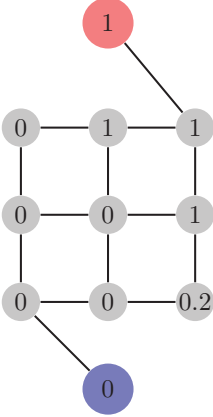
3.3.1. CCMF and CTV are not dual. Strang proved that the continuous max-flow problem is dual to the TV problem under some assumptions on g . Remarkably, this duality is not observed in the discrete case in which the CCMF dual and CTV problems are given by

$$\begin{aligned} \min_{\lambda, \nu} \quad & \lambda^T g^2 + \frac{1}{4} \left(\mathbf{1}^n \cdot / (|A|\lambda) \right)^T \left((c + A\nu)^2 \right) \quad \neq \quad \min_u g^T \sqrt{|A^T|(Au)^2} \\ \text{s.t.} \quad & \lambda \geq 0, \quad \text{s.t.} \quad u_s - u_t = 1. \end{aligned}$$

We note that the analytic expressions are different and also not equivalent. The duality of the classical max-flow problem with the minimum cut holds because in the expression of the Lagrangian function, it is possible to deduce a value of λ by substituting it into the dual problem so that it depends only on ν . However, λ in the dual CCMF problem (2.9) depends on several values of neighboring λ and ν . Thus, the CCMF dual problem cannot be

Table 1

Example illustrating the difference between optimal solutions of CCMF problem and CTV.

<p>Transport graph with weights g on the nodes</p>	<p>Combinatorial Continuous Max-Flow (CCMF)</p> $\max_F F_{st}$ <p>s.t. $A^T F = 0,$ $A^T F ^2 \leq g^2.$</p>	<p>Combinatorial Total Variation (CTV)</p> $\min_u g^T \sqrt{ A^T } (Au)^2$ <p>s.t. $u_s - u_t = 1.$</p>
		
	<p>$F_{st} = 1.63$</p>	$CTV(u) = \sum_i g_i \sqrt{\sum_j (u_i - u_j)^2}$ $= 2 + 1 \times \sqrt{2} + 1 \times \sqrt{2}$ $+ 1 \times \sqrt{1 + 0.8^2}$ $+ 2 \times \sqrt{0.8^2 + 0.2^2}$ $+ 2 \times 0.2$ $= 8.16.$

simplified by removing a variable, for example by identifying ν and u in the CTV problem. Said differently, combining a null-divergence constraint with an *anisotropic* capacity constraint in the classical max-flow formulation allows the duality to hold; in contrast, in the CCMF formulation, combining a null-divergence constraint with an *isotropic* capacity constraint does not allow such a duality to hold.

Numerically we can also show in examples that the value $\max F_{st}$ of the flow optimizing CCMF is not equal to the minimum value of CTV (See Table 1).

3.3.2. Theoretical links between CCMF and CTV. Even if CCMF is not dual with CTV, the two problems are weakly dual.

In the combinatorial setting the weak duality property is given by

$$(3.7) \quad \|F\| \leq g \Rightarrow F^T(Au) \leq g^T \|Au\|.$$

The next property shows that the norm $\|F\| = |A^T|F^2$ verifies weak duality.

Proposition 3.1. *Let G be a transport graph, and let F be a flow in G verifying the capacity constraint $\sqrt{|A^T|F^2} \leq g$. Let u be a vector of \mathbb{R}^n (defined on nodes of G). Then*

$$F^T Au \leq g^T \sqrt{|A^T|(Au)^2}.$$

Proof. Since we know that $\sqrt{|A^T|F^2} \leq g$, the following statement is true:

$$\sqrt{|A^T|F^2}^T \sqrt{|A^T|(Au)^2} \leq g^T \sqrt{|A^T|(Au)^2}.$$

We can now show that

$$F^T Au \leq \sqrt{(|A^T|F^2)^T} \sqrt{|A^T|(Au)^2}.$$

Using summation notation, then by the Cauchy–Schwarz inequality,

$$\sum_{i \in V} \sum_{e_{ij} \in E} F_{ij}(u_i - u_j) \leq \sum_{i \in V} \sqrt{\left(\sum_{e_{ij} \in E} F_{ij}^2 \right)} \left(\sum_{e_{ij} \in E} (u_i - u_j)^2 \right).$$

We conclude that

$$F^T Au \leq \sqrt{(|A^T|F^2)^T} \sqrt{|A^T|(Au)^2} \leq g^T \sqrt{|A^T|(Au)^2}. \quad \blacksquare$$

In terms of energy value, this proposition means that the CCMF energy, that is to say, the flow, is always smaller than the CTV.

Proposition 3.2. *Let F be a compatible flow verifying the constraints in (2.4), and let u be a vector of \mathbb{R}^n such as $u_s - u_t = 1$. Then,*

$$(3.8) \quad F_{st} \leq g^T \sqrt{|A^T|(Au)^2}.$$

Proof. In the combinatorial setting, the Green formula gives

$$F^T Au = u^T A^T F.$$

Let a be a vector of \mathbb{R}^n defined for each node v_i as

$$(3.9) \quad a_i = \begin{cases} -1 & \text{if } v_i \text{ belongs to the sink,} \\ 1 & \text{if } v_i \text{ belongs to the source,} \\ 0 & \text{otherwise.} \end{cases}$$

We can provide the following equivalent formulation of the CCMF problem (2.4) using an incidence matrix A of the transport graph deprived of the sink-source edge:

$$(3.10) \quad \begin{aligned} & \max F_{st} \\ \text{s.t. } & A^T F = F_{st} a, \\ & |A^T|F^2 \leq g. \end{aligned}$$

As F verifies the divergence-free constraint $A^T F = F_{st} a$,

$$u^T A^T F = u^T F_{st} a,$$

and as the equation $u_s - u_t = 1$ may be written equivalently as $a^T u = 1$, we conclude that

$$u^T F_{st} a = F_{st},$$

and by weak duality (Property 3.1),

$$F_{st} \leq g^T \sqrt{A^T (A u)^2}. \quad \blacksquare$$

This property should be of interest for extending CCMF to applications other than segmentation, which is outside the scope of this paper. In the next section, we present the performance of our approach in the context of segmentation.

4. Results. We now present applications of CCMF in image segmentation. To be used as a segmentation algorithm, three solutions are possible: the dual variable ν may be used directly if the user needs a matted result; otherwise ν may be thresholded; or finally an isoline or isosurface may be extracted from ν .

In the introduction we discussed how several works are related to CMF. Some are equivalent or slight modifications of AT-CMF for nonsegmentation applications [46, 45, 38, 48]. As in the original AT-CMF work, none of these comes with a convergence proof. In contrast, the work of [37], [32] are provably convergent, but both are very slow and do not generalize easily to 3D image segmentation. Consequently, it seems reasonable to compare our segmentation method to the original AT-CMF as representative of the continuous approach, as well as graph cuts, which represent the purely discrete case. Finally, even if the TV problem is different from the CMF problem (existence of a duality gap in the continuous setting [35]), the two problems are related enough to be compared. Specifically, we will also compare with Chambolle and Pock's recent work [11], which presented an algorithm for optimizing a TV-based energy for image segmentation.

Our validation is intended to establish three properties of the CCMF algorithm. First, we establish that CCMF does avoid the metrication artifacts exhibited by conventional graph cuts (on a 4-connected lattice). This property is established by examples on a natural image and the recovery of the classical catenoid structure as the minimal surface spanning two rings. Second, we compare the convergence of the CCMF algorithm to the AT-CMF algorithm to show that the CCMF algorithm converges more quickly and in a more stable fashion. Finally, we establish that our formulation of the CMF problem does not degrade segmentation performance on a standard database. In fact, because of the reduction in metrication error our algorithm gives improved numerical results. For this experiment, we use the GrabCut database to compare the quality of the segmentation algorithms. In addition to the above tests, we demonstrate through examples that the CCMF algorithm is also flexible enough to incorporate prior (unary) terms and to operate in three dimensions.

4.1. Metrication artifacts and minimal surfaces. We begin by comparing the CCMF segmentation result with the classical max-flow algorithm (graph cuts). Figure 3 shows the




Figure 3. *Brain segmentation.* (a) Original image with foreground and background seeds. (b), (c), (d) Segmentation obtained with (b) graph cuts (GC), (c) AT-CMF, threshold of P (obtained after 10000 iterations), (d) CCMF, threshold of ν (15 iterations).

segmentation of a brain, in which the contours obtained by graph cuts are noticeably blocky in the areas of weak gradient, while the contours obtained by both AT-CMF and CCMF are smooth.

In the continuous setting, the maximum flow computed in a 3D volume produces a minimal surface. The CCMF formulation may be also recognized as a minimal surface problem. In the dual formulation, the objective function is equivalent to a weighted sum of surface nodes. In [1], Appleton and Talbot compared the surfaces obtained from their algorithm with the analytic solution of the catenoid problem to demonstrate that their algorithm was a good approximation of the continuous minimal surface and was not creating discretization artifacts. The catenoid problem arises from consideration of two circles with equal radius whose centers lie along the z axis. The minimal surface which forms to connect the two circles is known as a catenoid. The catenoid appears in nature, for example, by creating a soap bubble between two rings. In order to demonstrate that CCMF is also finding a minimal surface, we performed the same catenoid experiment as that in [1]. The results are displayed in Figure 4, where we show that CCMF approximates the analytical solution of the catenoid with even greater fidelity than the AT-CMF example.

In order to verify that the solution obtained by CCMF approximates perimeters of planar objects in two dimensions, we segmented several shapes—squares, discs—with known perimeters, scaled at different sizes, and compared the analytic perimeters with the energies obtained by CCMF. The results are reported in Figure 5, where we observe that the true boundary length is proportional to the energy obtained with CCMF.

4.2. Stability, convergence, and speed. We may compare the segmentation results using ν to Appleton and Talbot’s result using P . We recall that AT-CMF solves the PDE system (2.3) in order to solve the continuous maximum flow problem (2.1), but no proof was given for convergence. The potential function P approximates an indicator function, with 0 values for




Figure 4. The catenoid test problem: The source is constituted by two full circles and sink by the remaining boundary of the image. (a) Surface computed analytically, (b) isosurface of P obtained by CMF, (c) isosurface of ν obtained by CCMF. The root mean square error (RSME) has been computed to evaluate the precision of the results to the surface computed analytically. The RSME for CMF is 1.98 and for CCMF is 0.75. The difference between those results is due to the fact that the CMF algorithm enforces exactly the source and sink points, leading to discretization around the disks. In contrast, the boundary localized around the seeds of ν is smooth, composed of grey levels. Thus the resulting isosurface computed by CCMF is more precise.

the background labels and 1 for the foreground. It can be difficult to know when to stop the AT-CMF algorithm, since the iterations used to solve the AT-CMF algorithm may oscillate, as displayed in Figure 6 on a synthetic image. In contrast, the CCMF algorithm is guaranteed to converge and smoothly approaches the optimum solution.

We also compare segmentation results with results obtained by the recent algorithm of Chambolle and Pock [11] optimizing a TV-based energy using a dual variable. In comparison to our work, the dual variable is not a flow (no constraint on the divergence and no proof of convergence toward a minimal surface as defined in [43] or [1]). In the literature, the continuous-domain duality between TV and max-flow is assumed, but in computational practice they are really different. Nevertheless, even as the method of Chambolle and Pock is solving a different problem, we performed numerical comparative tests. The problem of Chambolle and Pock (now denoted CP-TV) for binary segmentation optimizes

$$(4.1) \quad \min_{u \in [0,1]^n} \max_{F \in \mathbb{R}^m, \|F\|_\infty \leq 1} F^T(Au) + g^T u + \text{hard constraint attachment},$$

where $A \in \mathbb{R}^{m \times n}$ is the incidence matrix of the graph of n nodes and m edges defined on the image, and $g \in \mathbb{R}^n$ is the metric on the nodes, defined in (2.6).

The CCMF algorithm is faster than both AT-CMF and CP-TV for 2D image segmentation. We have implemented CCMF in MATLAB. We used an implementation of the Appleton–Talbot algorithm in C++ provided by the authors, and a MATLAB software implementing CP-TV on CPU, also provided by the respective authors. The three implementations make use of multithreaded parallelization, and the CPU times reported here were computed on an Intel Core 2 Duo (CPU 3.00GHz) processor, with 2 Gb of RAM. The CCMF average computation time on a 321×481 image is 181 seconds after 21 iterations. For AT-CMF, 80000 iterations require 547 seconds. For CP-TV, the average number of iterations needed for convergence is 36000, requiring an average time of 1961 seconds on CPU, and the median computation time is 1406 seconds. Note that a GPU implementation of CP-TV, also provided by the authors, achieves a speed gain factor of about 100 on an NVidia Quadro 3700. Although

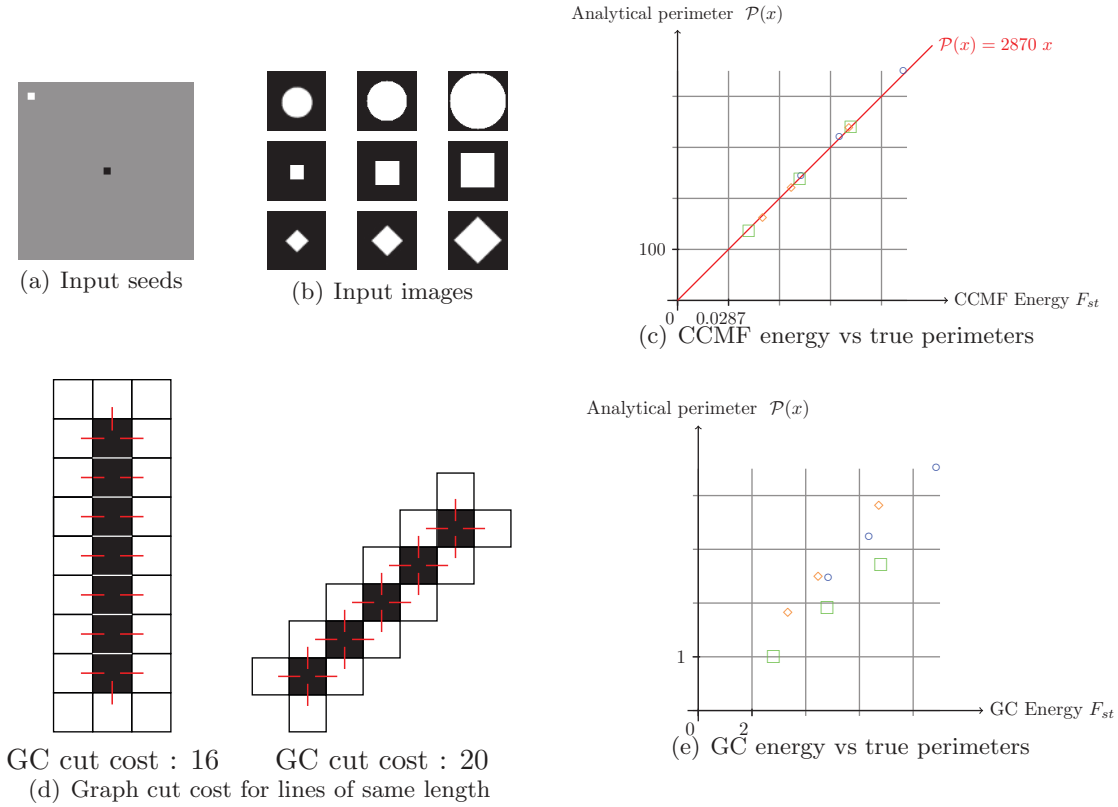


Figure 5. Comparison of perimeters computed analytically and with graph cuts and CCMF. To estimate the boundary length with CCMF, segmentations have been produced using the input images in (b) with the foreground and background seeds that appear in (a). (c) Plot showing the linear relation between the true perimeters and the energies obtained with CCMF. The symbols used to represent the dots correspond to the objects segmented in the images. A similar relation of proportionality is obtained with CP-TV. (d) Value of the cut obtained for a diagonal and a vertical line of same length using graph cuts. (e) Plot showing the nonlinear relation between the true perimeters and the energies obtained with graph cuts.

it was not pursued here, it should be noted that the CCMF optimization approach could also fully benefit from parallelization (e.g., on a GPU) because the core computation is to solve a linear system of equations. Realizing the benefit of a GPU solution would require the use of an iterative algorithm such as conjugate gradients or multigrid [3, 31], which would make comparisons to the direct solver used here less immediate. However, previous examples in the literature have suggested that this change in solver has not created difficulties in order to realize the benefits of a GPU architecture [23].

Sometimes, it may not always be necessary to wait so long for the complete convergence of CCMF (or AT-CMF) to obtain acceptable results. In cases where images exhibit sufficiently strong gradients, one iteration may be enough to obtain a satisfying segmentation. On such an image, one iteration of CCMF and 100 iterations of AT-CMF show acceptable approximate results reached only after about 2 seconds for either algorithm. However, one cannot always rely on strong edges, and the power of our CCMF formulation is to behave in a predictable fashion even when edges are weak.




Figure 6. Segmentation of an artificial image with AT-CMF (top row) and CCMF (bottom row). Top row (AT-CMF): (a) Image where the black and white discs are seeds. AT-CMF result stopped in (b) after 100 iterations, (c) 1000 iterations, (d) 10000 iterations. Bottom row (CCMF): (e) Image where the black and white discs are seeds. CCMF result (f) ν after 1 iteration, (g) ν after 15 iterations, and (h) threshold of the final ν .

We may also compare the computation time of CCMF to the optimization of TV using the split Bregman method [21]. Optimizing TV on a 100×100 image requires 5000 iterations and takes 23 seconds with a sequential implementation of split Bregman in C. On the same image, CCMF required only 17 iterations to reach the convergence criteria $\|r_d\| < 1$ and $\|\hat{\eta}\| < 2$, taking 4.7 seconds with a MATLAB implementation (using a 2-threaded solver). We conclude that TV-based methods appear to be quite slow in the context of image segmentation, although we employed the split Bregman algorithm, which is known as one of the fastest algorithms for TV optimization for image denoising, and the very recent CP-TV algorithm. This difference in speed between the denoising and segmentation applications is due to the very large number of iterations required for convergence to a binary segmentation.

4.3. Segmentation quality. In this experiment, we compare our CCMF formulation to the AT-CMF formulation and conventional graph cuts on the problem of image segmentation, to determine if there are strong performance differences between the formulations. We expect that there will not be, since in principle all three formulations are trying to “minimize the cut” but have slightly different definitions of the cut length. The primary advantage that we expect from our formulation is in the reduction of metrication artifacts (as compared to conventional graph cuts) and speed/convergence (as compared to AT-CMF).

Our experiment consists of testing graph cut, AT-CMF, CP-TV, and CCMF algorithms on a database with the same seeds. We used the Microsoft “GrabCut” database available online [39], which is composed of 50 images provided with seeds. From the seed images provided,




Figure 7. (a) Seed images belonging to the first (top) and second (bottom) set of seeds provided by the GrabCut database. (b) Segmentations by graph cuts computed with seeds from the two sets. (c) Segmentations by CCMF.

it is possible to extract two different possible markers for the background seed. Examples of such seeds are shown in Figure 7(a). Different convergence criteria are available for CCMF, such as the duality gap and norms of the residuals. However, for the CMF algorithm, we do not have any satisfying criteria. Bounding the number of nonbinary occurrences of P does not mean that the convergence is reached, because of possible oscillations. An intermediate result after 10000 iterations may be significantly different from the result reached after 1000000. Consequently, we ran the AT-CMF algorithm until we were convinced of having reached convergence, i.e., when P was nearly binary and did not change significantly when we doubled the number of iterations. For half of the images in the GrabCut database, AT-CMF algorithm required more than 20000 iterations to reach convergence; for a third of the images, more than 80000 iterations; and for 1/4 of the images, more than 160000 iterations. Binary convergence was still not reached after even 500000 iterations for the rest of the images (1/4), but we stopped the computation anyway. The Chambolle–Pock TV-based algorithm for binary segmentation is provably convergent and benefits also from several possible convergence criteria. The convergence criterion used is a ratio of changed pixels from two successive intervals of iterations being smaller than an epsilon that we fixed to 10^{-7} in order to obtain satisfying results. In practice, 1/4 of the images converged in less than 60000 iterations, and for the rest of them we increased the maximum number of iterations up to 200000. In contrast, CCMF needed only 21 iterations on average, and never more than 27, to reach the convergence criteria $\|r_d\| < 1$ and $\|\hat{\eta}\| < 2$. We noted that for all methods, segmentation quality degraded quickly if these conditions were not respected.

Table 2 displays the performance results for these algorithms. The Dice coefficient is a similarity measure between sets (segmentation and ground truth), ranging from 0 to 1.0 for

Table 2

Dice coefficient (percentage) computed between the segmentation mask and the ground truth image (provided by GrabCut database). The lines above the double bar show results for the first set of seeds, and the lines below the double bar show results obtained with the second set of seeds.

	Dice coeff.	GC	AT-CMF	CP-TV	CCMF
First set of seeds	mean	95.2	94.9	95.2	95.3
	median	96.2	96.1	96.3	96.3
	stand. dev.	4.3	4.1	4.1	4.1
Second set of seeds	mean	89.3	89.7	89.1	89.5
	median	91.2	92.0	91.7	92.3
	stand. dev.	8.4	7.7	8.5	8.6

no match and a perfect match, respectively. All the tested algorithms show very good results, with a Dice coefficient of 0.95–0.97 for the well positioned seeds, and 0.89–0.92 for the second set of seeds, further away from the objects. The CCMF and AT-CMF results are really close, and the mean is better than the GC and the CP-TV results.

4.4. Extensions. The primary focus of our experiments were to demonstrate that our CCMF algorithm achieves a solution which does not exhibit metrication artifacts, that it is fast with provable convergence, and that the segmentation quality is high. In the remainder of this section, we show that the algorithm can also incorporate unary terms and be equally applied in three dimensions. In fact, since CCMF is defined for an arbitrary graph, it could even be used to perform clustering in this more generalized framework. However, the benefit of avoiding gridding artifacts is less meaningful when performing clustering on an arbitrary graph, and therefore we omit any examples of this nature.

4.4.1. Unary terms. A simple modification of the transport graph G permits the use of unary terms to automatically specify objects to be segmented. This approach is similar to the use of unary terms in the graph cut computation [5]. The placement of unary terms for adding data attachment constraints in the max-flow problem is performed by adding edges linking every node of the lattice to the source and to the sink. In the case of the CCMF problem, as weights are defined on the nodes, we need to add intermediary nodes between the grid and source on the one hand, and the grid and sink on the other. However, since these intermediary nodes have just two edges incident on them, these node weights are equivalent to edge weights, meaning that our construction is equivalent to the use of unary terms for AT-CMF that was pursued by Unger, Pock, and Bischof [45]. In our case, we used weighted intermediary nodes to keep the CCMF framework consistent. Considering that the original lattice is composed of n nodes, we add for each node an “upper” node linked to the source and a “lower” node linked to the sink. An example of this construction is shown in Figure 8.

The weights of the additional nodes may be set to reflect the node priors for a particular application. For image segmentation, given mean background and foreground colors BC and FC , we can set the capacities of the background nodes to $g_i = \exp(-\beta(BC_i - I_i)^2)$ and the foreground nodes to $g_i = \exp(-\beta(FC_i - I_i)^2)$. Examples of results using these appearance priors are shown in Figure 8.

4.4.2. Classification. As the general CCMF formulation is defined on arbitrary graphs, it is possible to employ it in tasks beyond image segmentation. For instance, CCMF may be



Figure 8. Example of graph construction to incorporate unary terms for unsupervised segmentation. Each node is connected to a source node and sink node through an intermediary node which is weighted to reflect the strength of the positive and negative unary terms. This construction is then applied with a simple appearance prior to automatically segment two images.

employed to solve general problems of graph clustering, even on networks with no meaningful embedding such as social networks. We show a possible application in a classification example, considering the social network studied by Zachary [47]. After the split of a university karate club due to a conflict between its two leaders, Zachary’s goal was to study if it was possible to predict the side joined by the members, based only on the social structure of the club. Classically, the graph is built by associating each member with a node, and edges link two members when special affinities are known between them. As weights are associated here with the strength of coupling between members, different strategies are possible for assigning the weights onto the nodes, which is necessary for using CCMF. For example, a given node/member may be assigned by the mean of affinity with its neighboring members in the graph. This weighting strategy works well in this example (see Figure 9). Another strategy would be to compute the solution in the dual graph, that is, in the graph where nodes are replaced by edges and (weighted) edges are replaced by (weighted) nodes. We have shown in an example the ability of CCMF for classification, a task that cannot be treated with finite–element–based methods such as AT-CMF. Evaluating the performances of CCMF for classification may be a topic for future research.

4.4.3. 3D segmentation. For 3D image segmentation, the minimal surface properties of CCMF generate good quality results, as shown in Figure 10. The CCMF formulation applies equally well in two or three dimensions, since CCMF is formulated on an arbitrary graph (which may be a 2D lattice, a 3D lattice, or an even more abstract graph). In three dimensions, our CCMF implementation is suffering from memory limitations in the direct solver we used, limiting its performances. Future work will address this issue using a dedicated solver.

5. Conclusion. In this paper, we have presented a new combinatorial formulation of the continuous maximum flow problem and a solution using an interior point convex primal-

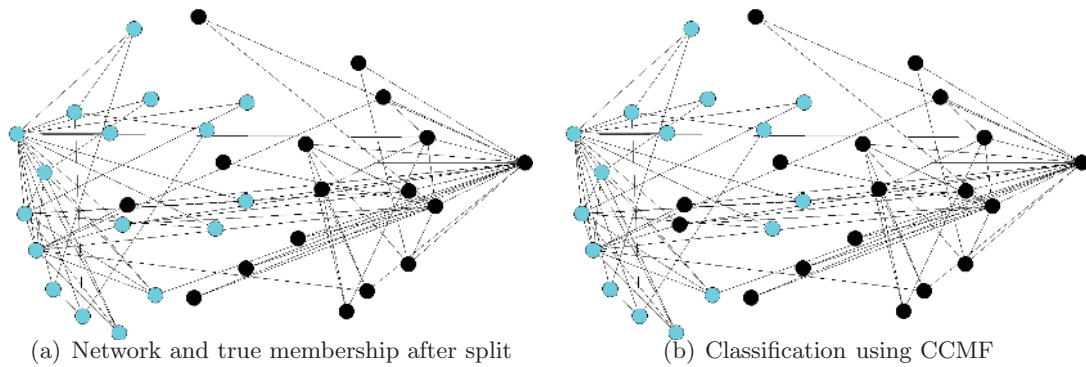


Figure 9. Zachary's karate club network. The two leaders in conflict are represented by the left-hand side and right-hand side nodes. Just one node is misclassified with CCMF. However, this node is known classically to be an unusual situation within the social network (see [47]) which is not captured by any known algorithm.




Figure 10. Liver segmentation in three dimensions by CCMF.

dual algorithm. This formulation allows us to optimize the maximum flow problem as well as its dual for which we provide an interpretation as a minimal surface problem. This new combinatorial expression of continuous max-flow avoids blockiness artifacts compared to graph cuts. Furthermore, the formulation of CMF on a graph reveals that it is actually the fact that the capacity constraints are applied to pointwise norms of flow *through nodes* that allows us to avoid metrication errors, as opposed to the conventional graph cut capacity constraint *through individual edges*. Additionally, it was shown that our CCMF formulation provides a better approximation to the analytic catenoid than the conventional AT-CMF discretization of the continuous max-flow problem. Contrary to graph cuts, we showed that CCMF estimates the true boundary length of objects. Finally, unlike finite-element-based methods such as AT-CMF, the CCMF formulation is expressed for arbitrary graphs and thus can be employed in a large variety of tasks such as classification.

We provide in this paper an exact analytic expression of the dual problem, convergence of the algorithm being guaranteed by the convexity of the problem. In terms of speed, when an approximate solution is sufficient, our implementation of CCMF in MATLAB is competitive with the Appleton–Talbot approach, which uses a system of PDEs, and a C++ implementation. The Appleton–Talbot algorithm has the significant drawback of not providing a criterion for convergence. In practice, this translates to long computation times when convergence is difficult for the AT-CMF algorithm. In contrast, our algorithm in this case is much faster, as we observe that the number of iterations required for convergence does not vary much for CCMF (less than a factor of two). The CCMF algorithm is simple to implement and may be applied to arbitrary graphs. Furthermore, it is straightforward to add unary terms to perform unsupervised segmentation.

We have also compared our algorithm to known weighted combinatorial TV minimization methods (CTV optimized with split Bregman, CP-TV). We have shown that our results are generally better for segmentation than combinatorial TV-based methods, and that our implementation on CPU is much faster and more predictable. Specifically, the number of iterations required for convergence does not vary much for CCMF, whereas it can vary by more than a factor of 10 for CP-TV. The deep study of the relationships between CCMF and combinatorial TV reveals that, in contrast with expectations from their duality under restrictive assumptions in the continuous domain, their duality relationship is weak only in the combinatorial setting. In fact, max-flow and TV are different problems with different constraints, yielding different algorithms and different results. One key difference between max-flow and TV is that max-flow algorithms were developed as segmentation algorithms, following the graph cut framework. One consequence is that max-flow formulations have a null-divergence objective, which is not present in TV formulations. Null divergence can be viewed as a consequence or a necessity in order to obtain constant partitions almost everywhere, in particular binary ones. This difference is important because in the proposed CCMF framework we *impose* a tight null-divergence constraint, which to our knowledge is novel for an isotropic formulation (graph-cuts have always had this constraint). AT-CMF and similar frameworks achieve null divergence if and when convergence is achieved. There is no null-divergence constraint at all in TV frameworks. As a consequence, while in practice it is possible to compare TV and max-flow formulations, they should be treated differently. However, the strong computational performance and segmentation quality results of CCMF as compared to TV (using the strongest known optimization methods) suggests that it may be advantageous to apply our CCMF formulation to problems for which TV has proven effective (such as filtering).

Several further optimizations of CCMFs are possible, for instance, using multigrid implementations, the possibility of using GPU to solve the iterative linear system, and using a dedicated solver for the particular sparse linear system involved in the computation. We also plan to compare the efficiency obtained with the interior point method to a first order algorithm for solving the CCMF problem.

Ultimately, we hope to employ CCMF as a powerful segmentation algorithm which avoids metrication artifacts and provides a fast solution with provable convergence. Furthermore, we intend to explore the potential of CCMF to optimize other energy functions for which graph cuts or TV have proven useful, such as surface reconstruction or efficient convex filtering.

REFERENCES

- [1] B. APPLETON AND H. TALBOT, *Globally minimal surfaces by continuous maximal flows*, IEEE Trans. Pattern Anal. Mach. Intell., 28 (2006), pp. 106–118.
- [2] A. BHUSNURMATH AND C. J. TAYLOR, *Graph cuts via ℓ_1 norm minimization*, IEEE Trans. Pattern Anal. Mach. Intell., 30 (2008), pp. 1866–1871.
- [3] J. BOLZ, I. FARMER, E. GRINSPUN, AND P. SCHRÖDER, *Sparse matrix solvers on the GPU: Conjugate gradients and multigrid*, ACM Trans. Graphics, 22 (2003), pp. 917–924.
- [4] S. BOYD AND L. VANDENBERGHE, *Convex Optimization*, Cambridge University Press, Cambridge, UK, 2004.
- [5] Y. BOYKOV AND M. P. JOLLY, *Interactive graph cuts for optimal boundary and region segmentation of objects in n -d images*, in Proceedings of ICCV'01, Vol. 1, 2001, pp. 105–112.
- [6] Y. BOYKOV AND V. KOLMOGOROV, *Computing geodesics and minimal surfaces via graph cuts*, in Proceedings of ICCV'03, 2003, pp. 26–33.
- [7] Y. BOYKOV, O. VEKSLER, AND R. ZABIH, *Fast approximate energy minimization via graph cuts*, IEEE Trans. Pattern Anal. Mach. Intell., 23 (2001), pp. 1222–1239.
- [8] A. BUADES, B. COLL, AND J. M. MOREL, *A non-local algorithm for image denoising*, in Proceedings of CVPR 2005, Vol. 2, 2005, pp. 60–65.
- [9] V. CASELLES, R. KIMMEL, AND G. SAPIRO, *Geodesic active contours*, Int. J. Comput. Vis., 22 (1997), pp. 61–79.
- [10] A. CHAMBOLLE, D. CREMERS, AND T. POCK, *A Convex Approach for Computing Minimal Partitions*, Technical report 649, Ecole Polytechnique CMAP, Palaiseau Cedex, France, 2008.
- [11] A. CHAMBOLLE AND T. POCK, *A first-order primal-dual algorithm for convex problems with applications to imaging*, J. Math. Imaging Vision, 40 (2011), pp. 120–145.
- [12] A. CHAMBOLLE, *An algorithm for total variation minimization and applications*, J. Math. Imaging Vision, 20 (2004), pp. 89–97.
- [13] T. F. CHAN, S. ESEDOGLU, AND M. NIKOLOVA, *Algorithms for finding global minimizers of image segmentation and denoising models*, SIAM J. Appl. Math., 66 (2006), pp. 1632–1648.
- [14] J. DARBON AND M. SIGELLE, *Image restoration with discrete constrained total variation part ii: Levelable functions, convex priors and non-convex cases*, J. Math. Imaging Vision, 26 (2006), pp. 277–291.
- [15] A. ELMOATAZ, O. LÉZORAY, AND S. BOUGLEUX, *Nonlocal discrete regularization on weighted graphs: A framework for image and manifold processing*, IEEE Trans. Image Processing, 17 (2008), pp. 1047–1060.
- [16] C. L. FEFFERMAN, *Existence and Smoothness of the Navier–Stokes Equation*, http://www.claymath.org/millennium/Navier-Stokes_Equations/navierstokes.pdf.
- [17] L. R. FORD AND D. R. FULKERSON, *Maximal flow through a network*, Canad. J. Math., 8 (1956), pp. 399–404.
- [18] G. GILBOA AND S. OSHER, *Nonlocal linear image regularization and supervised segmentation*, Multiscale Model. Simul., 6 (2007), pp. 595–630.
- [19] A. V. GOLDBERG AND R. E. TARJAN, *A new approach to the maximum-flow problem*, J. Assoc. Comput. Mach., 35 (1998), pp. 921–940.
- [20] D. GOLDFARB AND W. YIN, *Second-order cone programming methods for total variation-based image restoration*, SIAM J. Sci. Comput., 27 (2005), pp. 622–645.
- [21] T. GOLDSTEIN AND S. OSHER, *The split Bregman method for L_1 -regularized problems*, SIAM J. Imaging Sci., 2 (2009), pp. 323–343.
- [22] L. GRADY AND C. ALVINO, *The piecewise smooth Mumford-Shah functional on an arbitrary graph*, IEEE Trans. Image Process., 18 (2009), pp. 2547–2561.
- [23] L. GRADY, T. SCHIWETZ, S. AHARON, AND R. WESTERMANN, *Random walks for interactive organ segmentation in two and three dimensions: Implementation and validation*, in Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention, 2005, pp. 773–780.
- [24] L. GRADY AND E. L. SCHWARTZ, *Faster graph-theoretic image processing via small-world and quadtree topologies*, in Proceedings of CVPR, Vol. 2, 2004, pp. 360–365.

- [25] L. J. GRADY AND J. R. POLIMENI, *Discrete Calculus: Applied Analysis on Graphs for Computational Science*, Springer, London, 2010.
- [26] A. N. HIRANI, *Discrete Exterior Calculus*, Ph.D. thesis, California Institute of Technology, Pasadena, CA, 2003.
- [27] M. IRI, *Theory of flows in continua as approximation to flows in networks*, in *Survey of Mathematical Programming*, North-Holland, Amsterdam, 1979, pp. 263–278.
- [28] H. ISHIKAWA, *Exact optimization for Markov random fields with convex priors*, *IEEE Trans. Pattern Anal. Mach. Intell.*, 25 (2003), pp. 1333–1336.
- [29] M. KASS, A. WITKIN, AND D. TERZOPOULOS, *Snakes: Active contour models*, *Int. J. Comput. Vis.*, 1 (1988), pp. 321–331.
- [30] V. KOLMOGOROV AND R. ZABIN, *What energy functions can be minimized via graph cuts?*, *IEEE Trans. Pattern Anal. Mach. Intell.*, 26 (2004), pp. 147–159.
- [31] J. KRUGER AND R. WESTERMANN, *Linear algebra operators for GPU implementation of numerical algorithms*, in *Proceedings of SIGGRAPH 2003*, 2003, pp. 908–916.
- [32] R. A. LIPPERT, *Discrete approximations to continuum optimal flow problems*, *Stud. Appl. Math.*, 117 (2006), pp. 321–333.
- [33] Y. NESTEROV, *Gradient Methods for Minimizing Composite Objective Function*, CORE discussion papers 2007076, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), Louvain-la-Neuve, Belgium, 2007.
- [34] R. NOZAWA, *Max-flow min-cut theorem in an anisotropic network*, *Osaka J. Math.*, 27 (1990), pp. 805–842.
- [35] R. NOZAWA, *Examples of max-flow and min-cut problems with duality gaps in continuous networks*, *Math. Programming*, 63 (1994), pp. 213–234.
- [36] N. PARAGIOS, *Geodesic Active Regions and Level Set Methods: Contributions and Applications in Artificial Vision*, Ph.D. thesis, INRIA Sophia Antipolis, Sophia Antipolis Cedex, France, 2000.
- [37] T. POCK, D. CREMERS, H. BISCHOF, AND A. CHAMBOLLE, *An algorithm for minimizing the Mumford-Shah functional*, in *Proceedings of ICCV'09*, 2009, pp. 1133–1140.
- [38] T. POCK, T. SCHOENEMANN, G. GRABER, H. BISCHOF, AND D. CREMERS, *A convex formulation of continuous multi-label problems*, in *Proceedings of ECCV'08*, 2008, pp. 792–805.
- [39] C. ROTHER, V. KOLMOGOROV, AND A. BLAKE, *“GrabCut”—Interactive foreground extraction using iterated graph cuts*, in *ACM Trans. Graphics*, 23 (2004), pp. 309–314.
- [40] L. I. RUDIN, S. OSHER, AND E. FATEMI, *Nonlinear total variation based noise removal algorithms*, *Phys. D*, 60 (1992), pp. 259–268.
- [41] J. A. SETHIAN, *Level Set Methods and Fast Marching Methods*, Cambridge University Press, Cambridge, UK, 1999.
- [42] J. H. D. SHULMAN AND J. Y. HERVÉ, *Regularization of discontinuous flow fields*, in *Proceedings of the Workshop on Visual Motion*, 1989, pp. 81–86.
- [43] G. STRANG, *Maximum flows through a domain*, *Math. Programming*, 26 (1983), pp. 123–143.
- [44] G. STRANG, *Maximum flows and minimum cuts in the plane*, *J. Global Optim.*, 47 (2010), pp. 527–535.
- [45] M. UNGER, T. POCK, AND H. BISCHOF, *Interactive Globally Optimal Image Segmentation*, Technical report ICG-TR-08/02, Graz University of Technology, Graz, Austria, 2008.
- [46] M. UNGER, T. POCK, W. TROBIN, D. CREMERS, AND H. BISCHOF, *TVSeg - Interactive total variation based image segmentation*, in *Proceedings of BMVC'08*, 2008.
- [47] W. ZACHARY, *An information flow model for conflict and fission in small groups*, *J. Anthropol. Res.*, 33 (1977), pp. 452–473.
- [48] C. ZACH, D. GALLUP, J.-M. FRAHM, AND M. NIETHAMMER, *Fast global labeling for real-time stereo using multiple plane sweeps*, in *Proceedings of VMV'08*, 2008, pp. 243–252.
- [49] X. ZHANG, M. BURGER, X. BRESSON, AND S. OSHER, *Bregmanized nonlocal regularization for deconvolution and sparse reconstruction*, *SIAM J. Imaging Sci.*, 3 (2010), pp. 253–276.