

In A. Heyden, F. Kahl, C. Olsson, M. Oskarsson, X.-C. Tai (Eds.):
Energy Minimization Methods in Computer Vision and Pattern Recognition.
Lecture Notes in Computer Science, Vol. 8081, pp. 26–39, Springer, Berlin, 2013.
The final publication is available at link.springer.com.

Linear Osmosis Models for Visual Computing

Joachim Weickert¹, Kai Hagenburg¹, Michael Breuß², and Oliver Vogel¹

¹ Mathematical Image Analysis Group
Dept. of Mathematics and Computer Science, Campus E1.1
Saarland University, 66041 Saarbrücken, Germany
{weickert, hagenburg, vogel}@mia.uni-saarland.de

² Institute for Applied Mathematics and Scientific Computing
Platz der Deutschen Einheit 1, Brandenburg University of Technology
03046 Cottbus, Germany
breuss@tu-cottbus.de

Abstract. Osmosis is a transport phenomenon that is omnipresent in nature. It differs from diffusion by the fact that it allows nonconstant steady states. In our paper we lay the foundations of osmosis filtering for visual computing applications. We model filters with osmotic properties by means of linear drift-diffusion processes. They preserve the average grey value and the nonnegativity of the initial image. Most interestingly, however, we show how the nonconstant steady state of an osmosis evolution can be steered by its drift vector field. We interpret this behaviour as a data integration mechanism. In the integrable case, we characterise the steady state as a minimiser of a suitable energy functional. In the nonintegrable case, we can exploit osmosis as a framework to fuse incompatible data in a visually convincing way. Osmotic data fusion differs from gradient domain methods by its intrinsic invariance under multiplicative grey scale changes. The osmosis framework constitutes a novel class of methods that can be tailored to solve various problems in image processing, computer vision, and computer graphics. We demonstrate its versatility by proposing osmosis models for compact image representation, shadow removal, and seamless image cloning.

Keywords: osmosis, drift–diffusion, Fokker–Planck equation, diffusion filters, gradient domain methods, shadow removal, image editing

1 Introduction

While diffusion processes are frequently used in image processing, computer vision and computer graphics, there is a closely related transport phenomenon in

nature that is basically unexplored in visual computing applications: It is called osmosis [1]. Osmosis describes transport through a semipermeable membrane in such a way that in its steady state, the liquid concentrations on both sides of the membrane can differ. Osmosis is the primary mechanism for transporting water in and out of cells, and it has many applications in medicine and engineering. It can be seen as the nonsymmetric counterpart of diffusion. Since diffusion can only model symmetric transport processes, it leads to flat steady states [2].

Our Contributions. The goal of our paper is to lay the foundations of osmosis filtering for visual computing applications. In contrast to osmosis in natural systems we do not need two different phases (water and salt) and a membrane that is only permeable for one of them: We can obtain nonconstant steady states within a single phase that represents the grey value. All we have to do is to supplement diffusion with a drift term. The resulting drift-diffusion process is also in divergence form and thus preserves the average grey value of the initial image, but it allows to have full control over its nonflat steady state: We show that we can design osmosis filters that converge to any specified image. Most importantly, we shall see that osmosis has the ability to integrate conflicting gradient data in a visually convincing way. This enables many applications beyond classical data regularisation tasks. In particular, we show the potential of osmosis for three prototypical problems: compact data representation, shadow removal, and image editing. Interestingly, these applications do not require any nonlinearities: The richness of the drift term permits to reach these goals already within a linear setting.

Paper Structure. In Section 2 we describe our drift-diffusion framework for continuous osmosis filters, analyse its essential properties, and interpret osmosis processes as models for data integration. Afterwards we sketch a simple numerical scheme in Section 3. Applications of osmosis models to visual computing problems are described in Section 4, and our paper is concluded with a summary in Section 5.

Related Work. While diffusion filters are often combined with data fidelity terms, there are not many combinations with a drift term in divergence form so far. Hagenburg et al. [3] have proposed a lattice Boltzmann model for dithering that approximates a nonlinear drift-diffusion equation in the continuous limit. However, they did not investigate this continuous model any further. In [4] it was shown that a combination of a discrete osmosis model with a stabilised backward diffusion filter is useful for designing numerical schemes for hyperbolic conservation laws that benefit from low numerical diffusion. Moreover, this discrete osmosis model has been interpreted as a nonsymmetric Markov chain, while discrete diffusion filters lead to symmetric Markov chains. Illner and Neunzert [5] have investigated so-called directed diffusion processes that converge to a specified background image, but did not apply them to image processing problems.

With respect to their ability to integrate incompatible gradient data, osmosis methods can be related to gradient domain methods. In computer vision, gradi-

ent domain methods are used for shape from shading [6], for shadow removal [7], and as models for retinex [8]. In computer graphics they are useful for a number of image editing and fusion problems; see e.g. [9, 10]. Relations between osmosis and gradient domain methods are discussed in Section 2.3, and Section 4.3 gives an experimental comparison. With respect to their invariance under multiplicative brightness changes, osmosis methods also resemble Georgiev’s covariant derivative framework [11], but appear to be easier to understand.

Since our drift-diffusion formulation of osmosis filtering can be interpreted in a stochastic way as Fokker-Planck equation [12], it has some structural similarities to work by Sochen [13] that deals with a stochastic justification of the Beltrami flow. He mentions the potential benefit of drift terms but did not carry out any experiments. The Fokker-Planck equation has also been used by Wang and Hancock [14] for performing probabilistic relaxation labelling on graphs.

While the present paper focuses on the continuous theory and introduces specific models for different visual computing applications, we have also authored a companion paper that establishes a fully discrete theory for osmosis and studies efficient numerical methods [15].

2 Continuous Linear Osmosis Filtering

2.1 Drift–Diffusion Model

We consider a rectangular image domain $\Omega \subset \mathbb{R}^2$ with boundary $\partial\Omega$. A reasonable osmosis theory for greyscale images requires a positive initial image $f : \Omega \rightarrow \mathbb{R}_+$. Moreover, we assume that we can choose some *drift vector field* $\mathbf{d} : \Omega \rightarrow \mathbb{R}^2$. As we will see below, it allows us to steer the osmosis process to a desired nonflat steady state. Then a (*linear*) *osmosis filter* computes a family $\{u(\mathbf{x}, t) \mid t \geq 0\}$ of processed versions of $f(\mathbf{x})$ by solving the drift-diffusion equation

$$\partial_t u = \Delta u - \operatorname{div}(\mathbf{d}u) \quad \text{on } \Omega \times (0, T], \quad (1)$$

with f as initial condition,

$$u(\mathbf{x}, 0) = f(\mathbf{x}) \quad \text{on } \Omega, \quad (2)$$

and homogeneous Neumann boundary conditions:

$$\langle \nabla u - \mathbf{d}u, \mathbf{n} \rangle = 0 \quad \text{on } \partial\Omega \times (0, T]. \quad (3)$$

Here $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product, and \mathbf{n} is the outer normal vector to the image boundary $\partial\Omega$. Thus, the boundary conditions specify a vanishing flux across the image boundaries.

Extending linear osmosis to colour images does not create specific problems: One proceeds separately in each RGB channel using individual drift vector fields in each channel.

2.2 Theoretical Properties

While the main focus of our paper is on modelling aspects, successful modelling is impossible without some insights into essential theoretical properties. They are summarised in the following proposition.

Proposition 1. [Theory for Continuous Linear Osmosis]

A classical solution of the linear osmosis process (1)–(3) with positive initial image $f : \Omega \rightarrow \mathbb{R}_+$ and drift vector field $\mathbf{d} : \Omega \rightarrow \mathbb{R}^2$ satisfies the following properties:

(a) The average grey value is preserved:

$$\frac{1}{|\Omega|} \int_{\Omega} u(\mathbf{x}, t) \, d\mathbf{x} = \frac{1}{|\Omega|} \int_{\Omega} f(\mathbf{x}) \, d\mathbf{x} \quad \forall t > 0. \quad (4)$$

(b) The evolution preserves nonnegativity:

$$u(\mathbf{x}, t) \geq 0 \quad \forall \mathbf{x} \in \Omega, \quad \forall t > 0. \quad (5)$$

(c) If \mathbf{d} satisfies

$$\mathbf{d} = \nabla(\ln v) = \frac{\nabla v}{v} \quad (6)$$

with some positive image v , then the following holds:
The steady state equation

$$\Delta u - \operatorname{div}(\mathbf{d}u) = 0 \quad (7)$$

is equivalent to the Euler-Lagrange equation of the energy functional

$$E(u) = \int_{\Omega} v \left| \nabla \left(\frac{u}{v} \right) \right|^2 \, d\mathbf{x}. \quad (8)$$

Moreover, the steady state solution of the osmosis process is given by $w(\mathbf{x}) = \frac{\mu_f}{\mu_v} v(\mathbf{x})$, where μ_f and μ_v denote the average grey values of f and v .

Proof.

(a) Let $\mu(t) := \frac{1}{|\Omega|} \int_{\Omega} u(\mathbf{x}, t) \, d\mathbf{x}$ denote the average grey value at time $t \geq 0$. Using the divergence theorem and the homogeneous Neumann boundary conditions we obtain

$$\begin{aligned} \frac{d\mu}{dt} &= \frac{1}{|\Omega|} \int_{\Omega} \partial_t u \, d\mathbf{x} = \frac{1}{|\Omega|} \int_{\Omega} \operatorname{div}(\nabla u - \mathbf{d}u) \, d\mathbf{x} \\ &= \int_{\partial\Omega} \langle \nabla u - \mathbf{d}u, \mathbf{n} \rangle \, dS = 0. \end{aligned} \quad (9)$$

Thus, the average grey value remains constant over time.

- (b) Assume that $T > 0$ is the smallest time where $\min_{\mathbf{x}, t} u(\mathbf{x}, t) = 0$, and that this minimum is attained in some inner point $\boldsymbol{\xi}$. Thus, $\nabla u(\boldsymbol{\xi}, T) = \mathbf{0}$, and we have

$$\partial_t u(\boldsymbol{\xi}, T) = \Delta u(\boldsymbol{\xi}, T) - \underbrace{u(\boldsymbol{\xi}, T)}_{=0} \operatorname{div} \mathbf{d} - \mathbf{d}^\top \underbrace{\nabla u(\boldsymbol{\xi}, T)}_{=\mathbf{0}}. \quad (10)$$

This shows that in $(\boldsymbol{\xi}, T)^\top$ the osmosis evolution behaves like the diffusion equation $\partial_t u = \Delta u$. It is well known that for diffusion with homogeneous Neumann boundary conditions the minimum cannot decrease in time; see e.g. [2]. Thus, the solution of the osmosis process remains nonnegative.

- (c) The energy functional (8) can be rewritten as

$$E(u) = \int_{\Omega} F(u, \nabla u) \, d\mathbf{x} \quad (11)$$

with

$$F(u, \nabla u) = \frac{|v \nabla u - u \nabla v|^2}{v^3}. \quad (12)$$

From the calculus of variations we know that any minimiser of $E(u)$ satisfies the Euler-Lagrange equation

$$0 = F_u - \partial_x F_{u_x} - \partial_y F_{u_y} \quad (13)$$

with homogeneous Neumann boundary conditions, where $\mathbf{x} = (x, y)^\top$ and subscripts denote partial derivatives. With F from (12) this becomes after some simplifications

$$0 = -2v \operatorname{div} \left(\frac{v \nabla u - u \nabla v}{v^3} \right) - \frac{4 \nabla^\top v (v \nabla u - u \nabla v)}{v^3}. \quad (14)$$

Using

$$\operatorname{div} \left(v \nabla \left(\frac{u}{v} \right) \right) = v^2 \operatorname{div} \left(\frac{v \nabla u - u \nabla v}{v^3} \right) + \frac{2 \nabla^\top v (v \nabla u - u \nabla v)}{v^2} \quad (15)$$

the Euler-Lagrange equation can be written as

$$0 = -\frac{2}{v} \operatorname{div} \left(v \nabla \left(\frac{u}{v} \right) \right). \quad (16)$$

It is easy to check that this is equivalent to (7) if $\mathbf{d} = \frac{\nabla v}{v}$ with $v > 0$. Straightforward computations also show that one obtains (3) as boundary condition on $\partial\Omega$.

It is clear that an image v with $\mathbf{d} = \frac{\nabla v}{v}$ also fulfils the steady state equation (7) of the osmosis evolution with homogeneous Neumann boundary conditions. However, note that with v also cv with any constant c is a solution of this problem. Since the osmosis evolution preserves the average grey value and the nonnegativity of the initial image, it can only converge to a rescaled version w of v that is nonnegative and has the same average grey value as the original image f . Thus, $w(\mathbf{x}) = \frac{\mu_f}{\mu_v} v(\mathbf{x})$. \square

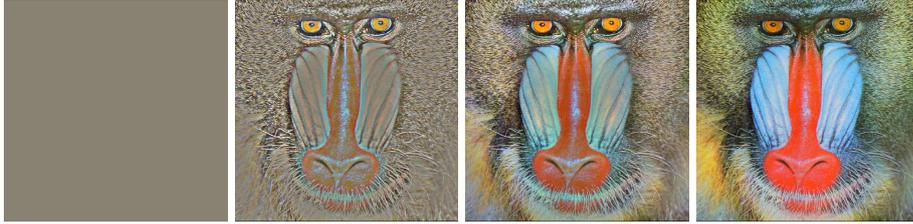


Fig. 1. Convergence of osmosis to a specified image. **From left to right:** (a) Original image, 512×512 pixels. Each channel has the same mean value as the *mandrill* test image. (b) Osmosis result at evolution time $t = 50$. (c) $t = 1000$. (d) $t = 250000$ gives a steady state that is identical to the mandrill image.

Preservation of the average grey value does not distinguish osmosis from diffusion filtering [2]. However, while diffusion filtering satisfies a maximum–minimum principle [2], osmosis only fulfils a weaker form of stability, namely preservation of nonnegativity. We conjecture that it is possible to establish preservation of strict positivity, since this also holds in the fully discrete case [15].

Proposition 1 implies that osmosis permits nontrivial steady states. This is a fundamental difference to diffusion that allows only flat steady states [2]. Of course, these steady state results should be accompanied by a formal convergence analysis. This is mathematically more involved and will be presented in a journal version of our paper.

Figure 1 illustrates such a convergence behaviour. Starting from a flat initial image, we can choose the drift vector field such that osmosis converges to the *mandrill* test image: If $\mathbf{v} = (v_1, v_2, v_3)^\top$ is the RGB image of the mandrill, the drift vector in channel i is chosen as $\nabla(\ln v_i)$.

2.3 Osmosis as a Process for Data Integration

We have seen that for $\mathbf{d} := \nabla(\ln v)$, osmosis converges to a multiplicatively rescaled version of v . This motivates us to call $\mathbf{d}[v] := \nabla(\ln v)$ the *canonical drift vector field* of the image v . Since $\mathbf{d}[v]$ contains derivative information of the steady state, we may regard osmosis as a process for data integration.

Obviously it is not very exciting to design an osmosis process that converges to an image which we know already. However, much more interesting situations arise when we modify the drift vector field, e.g. by setting certain components to zero, or by fusing the canonical drift vector fields of different images. Such applications will be considered in Section 4. Although in general the new drift vector field will be nonintegrable, osmosis will still create a steady state that aims at finding a good compromise between all conflicting constraints. In that sense osmosis resembles gradient domain methods that are popular both in computer

vision [6–8] and in computer graphics [10, 9]. Let us analyse these connections in more detail.

Gradient domain methods integrate a (possibly nonintegrable) gradient field approximation $\mathbf{p} = (p_1, p_2)^\top$ by minimising the energy

$$E(u) = \int_{\Omega} |\nabla u - \mathbf{p}|^2 d\mathbf{x}. \quad (17)$$

The corresponding Euler–Lagrange equation that a minimiser w has to fulfil is given by the Poisson equation

$$\Delta w = \operatorname{div} \mathbf{p}. \quad (18)$$

In the integrable case with $\mathbf{p} = \nabla v$, it is clear that an additive shift of v gives the same \mathbf{p} . Thus, gradient domain methods can recover v up to an additive constant.

On the other hand, an osmotic steady state w satisfies

$$\Delta w = \operatorname{div}(\mathbf{d}w). \quad (19)$$

An integrable osmosis setting with $\mathbf{d} = \nabla(\ln v)$ is invariant under multiplicative rescalings of v . In this sense it resembles Georgiev’s covariant derivative framework [11], but appear to be easier to comprehend. For computer vision applications where illumination changes are often modelled as multiplicative changes of the grey values, this multiplicative invariance of osmosis is preferable over the additive invariance of gradient domain methods.

In the nonintegrable case, gradient domain methods and osmosis can give different results that cannot be transformed into each other by simple additive or multiplicative grey value changes.

3 A Simple Numerical Scheme

To keep our paper self-contained, let us now sketch a simple explicit finite difference scheme for our osmosis model. For more numerical details and more efficient schemes we refer to [15], where a general fully discrete theory for osmosis filtering is established.

We consider a grid size h in x - and y -direction, and a time step size $\tau > 0$. Moreover, we denote by $u_{i,j}^k$ an approximation to u in the grid point $((i-\frac{1}{2})h, (j-\frac{1}{2})h)^\top$ at time $k\tau$. Setting $\mathbf{d} = (d_1, d_2)^\top$, a straightforward finite difference discretisation of (1) is given by

$$\begin{aligned} \frac{u_{i,j}^{k+1} - u_{i,j}^k}{\tau} &= \frac{u_{i+1,j}^k + u_{i-1,j}^k + u_{i,j+1}^k + u_{i,j-1}^k - 4u_{i,j}^k}{h^2} \\ &\quad - \frac{1}{h} \left(d_{1,i+\frac{1}{2},j} \frac{u_{i+1,j}^k + u_{i,j}^k}{2} - d_{1,i-\frac{1}{2},j} \frac{u_{i,j}^k + u_{i-1,j}^k}{2} \right) \\ &\quad - \frac{1}{h} \left(d_{2,i,j+\frac{1}{2}} \frac{u_{i,j+1}^k + u_{i,j}^k}{2} - d_{2,i,j-\frac{1}{2}} \frac{u_{i,j}^k + u_{i,j-1}^k}{2} \right). \quad (20) \end{aligned}$$

It allows to compute the results at time level $k + 1$ from the data at level k . This scheme also holds for boundary points, if we mirror the image at its boundaries and assume a zero drift vector across boundaries.

For some positive image v , we obtain a discrete approximation of its canonical drift vector field $(d_1[v], d_2[v])^\top = \frac{\nabla v}{v}$ at intermediate grid points via

$$d_{1,i+\frac{1}{2},j} = \frac{2(v_{i+1,j} - v_{i,j})}{h(v_{i+1,j} + v_{i,j})}, \quad d_{2,i,j+\frac{1}{2}} = \frac{2(v_{i,j+1} - v_{i,j})}{h(v_{i,j+1} + v_{i,j})}. \quad (21)$$

In [15] we show that the scheme (20)–(21) preserves positivity and converges to its unique steady state if the time step size satisfies

$$\tau < \frac{h^2}{8}. \quad (22)$$

4 Application to Visual Computing Problems

In order to illustrate the potential of osmosis models to solve visual computing problems, we study three fairly different applications: compact image representation, shadow removal, and seamless image cloning. All results below display steady states that have been computed with the numerical scheme from Section 3 with $h := 1$ and $\tau := 0.1$. One can achieve positivity of a bitwise coded initial image by adding an offset value of $\epsilon > 0$ such that each channel lies in the range $[\epsilon, 255 + \epsilon]$. Offset values should not be too large to avoid that they have a visible impact on the result. We choose $\epsilon := 1$.

4.1 Compact Data Representation

Let us now investigate if osmosis processes can be useful for compact image representations. There has been a long tradition of reconstructing images from their information near edges; see e.g. [16, 17]. One may for instance store the grey values on both sides of the edges as Dirichlet data, and interpolate the remaining data by solving the Laplace equation $\Delta u = 0$ in between. While this requires to store two grey values per edge point, it appears tempting to use osmosis and keep only the magnitude of the induced drift vector at each edge pixel, since we know that its direction is orthogonal to the edge contour. All drift vectors that are not adjacent to edges are set to zero, such that homogeneous diffusion interpolation is performed. This is illustrated in Figure 2. We observe that this compact image representation works well at step edges, while the contrasts that are reproduced at smooth edges appear somewhat too low. This proof-of-concept application indicates that osmosis can become a valuable tool for encoding step edges within a more comprehensive image compression approach based in partial differential equations (PDEs). More details will be reported in forthcoming publications.



Fig. 2. Osmosis for compact image representation. **(a) Left:** Original image. **(b) Middle:** Canny edges, amounting to 9.6% of all pixels. **(c) Right:** Reconstruction using the average grey value of the original image and the canonical drift vectors in the edges.

4.2 Shadow Removal

Many computer vision tasks such as segmentation, tracking, and object recognition benefit from the removal of shadows. To this end, one wants to identify the shadow region and adapt its brightness to the brightness of the rest of the image. Several methods have been proposed to find suitable shadow edges; see [18] and the references therein. Here we assume that the shadow edges are given, and we concentrate on the brightness adaptation problem. So far, this brightness adaptation has been achieved for example with gradient domain methods [7] or with pyramid-based approaches [18].

Interestingly the invariance of osmosis under global multiplicative greyscale changes offers a particularly elegant solution for this task: If one models shadows as a local multiplicative illumination change within the image, then this only affects the canonical drift vectors at the transition between the shadow and the rest of the image. Thus, shadow removal can be accomplished by simply modifying the canonical drift vectors at the shadow boundaries. Setting them to zero at these locations turns osmosis locally into homogeneous diffusion and guarantees a continuous transition. An osmosis evolution that starts with the original image and uses these modified drift vectors converges to a steady state where the shadow has been removed. Osmosis shadow removal has no problems recovering texture details in the shadow part which is an important criterion in state-of-the-art methods [19].

Figure 3 illustrates these ideas. If one uses the original image as initialisation, the results will be somewhat too dark due to the shadow region and the fact that osmosis preserves the average grey value (or colour value in each channel). As a remedy, one may want to rescale the results such that the values in the non-shadow region approximate the ones in the initial image. This is done in Figure 3(d).

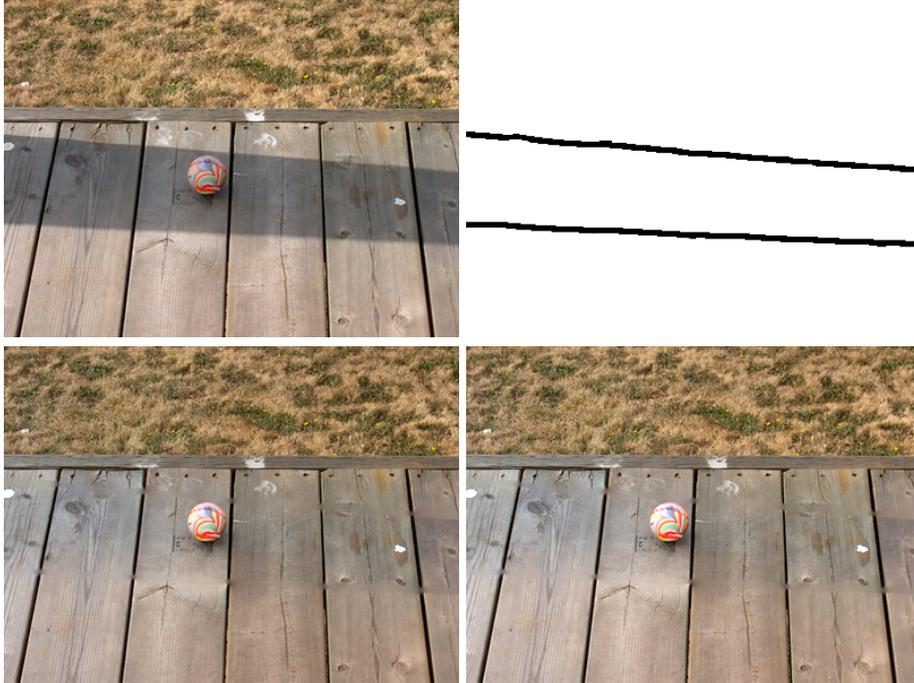


Fig. 3. Shadow removal by osmosis. **(a) Top left:** Original image (400×299 pixels). **(b) Top right:** User-selected shadow boundaries. In these boundaries all drift vectors are set to zero (homogeneous diffusion). In the other areas, the canonical drift vectors are used. **(c) Bottom left:** Osmosis reconstruction with (a) as initialisation. **(d) Bottom right:** Multiplicative rescaling of (c) such that the colours in the non-shadow regions approximate the ones in (a).

4.3 Seamless Image Cloning

The property of osmosis to fuse incompatible information can also be used for seamless image cloning. Figure 4 illustrates the problem: Two images f_1 and f_2 are to be merged such that f_2 replaces image information of f_1 . The rectangular image domain of the original image f_1 is denoted by Ω , and the image domain that is to be inserted is Γ . Its boundary is given by $\partial\Gamma$.

The classical gradient domain method for seamless image cloning is *Poisson image editing* [10]. It creates a fused image by solving the Poisson equation (18) with gradient data $\mathbf{p} = \nabla f_2$ in Γ and Dirichlet boundary conditions $u = f_1$ on $\partial\Gamma$. By construction, this localises the influence of the patch to the domain Γ .

To provide an osmosis-based alternative to Poisson image editing, we proceed as follows: We use the canonical drift vectors of f_1 in $\Omega \setminus \Gamma$, and the ones of f_2 in Γ . At the interface $\partial\Gamma$, we use the arithmetic mean of both drift vectors. The

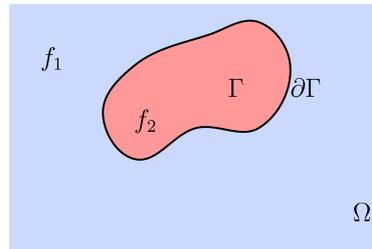


Fig. 4. Illustration of the image editing problem.



Fig. 5. Seamless image cloning (cf. [4]). **(a) Top left:** Painting of Euler. **(b) Top middle:** Drawing of Lagrange (with selected face region). **(c) Top right:** Direct cloning of Lagrange on top of Euler's face. **(d) Bottom left:** Poisson image editing with Dirichlet data of Euler at the interface. **(e) Bottom right:** Osmosis image editing with averaged drift vector fields at the interface. Source of original images: Wikimedia Commons.



Fig. 6. Comparison between Poisson image editing and osmosis image editing when the contrast of the inserted patch is not optimal. **(a) Top left:** Original image (400×300 pixels). **(b) Top right:** Direct cloning of a patch where the RGB values have been multiplied by a factor 0.25. **(c) Bottom left:** Poisson image editing suffers from the poor contrast within the cloned patch. **(d) Bottom right:** Osmosis image editing gives a much better contrast reconstruction.

process is initialised with f_1 on the entire rectangular image domain Ω , and its steady state gives the cloned image.

Figure 5 juxtaposes the results of Poisson image editing and osmosis image editing for an application where we want to clone the face of Lagrange on the body of Euler. While both methods give seamless results, Poisson image editing is unable to adapt the higher contrast of the face of Lagrange to the lower contrast of the Euler image. Section 2.3 gives an explanation for this phenomenon: Gradient domain methods allow only additive grey value shifts, and additive shifts cannot influence the contrast. Osmosis image editing, on the other hand, enables multiplicative changes that can also adapt the contrast. This is one reason why the osmosis result in Fig. 5(f) comes visually closer to the original Euler image in Fig. 5(a). A second reason results from the averaging of the canonical drift vectors at the interface $\partial\Gamma$. This makes the influence of the osmosis editing global, which also contributes to a more harmonic impression.

Fig. 6 presents a synthetic experiment that is tailored to visualise the differences between gradient domain editing and osmosis editing. We have reduced the contrast in the mouth region of the teddy bear by multiplying the RGB values by 0.25. While gradient domain editing cannot increase the contrast in the cloned patch, this is no problem for osmosis image editing.

These discussions show that there are two reasons for the superiority of osmosis image editing over Poisson image editing: boundary conditions that avoid a strict localisation, and the ability of osmosis to perform multiplicative brightness adaptations instead of additive ones.

5 Summary and Conclusions

We have advocated osmosis as a novel concept for visual computing. It is surprising that after decades of intensive research on PDE methods in image analysis, this important process in nature has been widely ignored by our research community so far. While osmosis filters differ from homogeneous diffusion filtering “only” by their drift term, this term has a fundamental consequence: It creates nonconstant steady states that can be controlled in a transparent way by the drift vector field. This offers interesting application areas that go far beyond the classical image regularisation and enhancement applications of diffusion methods. We have illustrated this potential by using osmosis for compact image representation, shadow removal, and seamless image cloning. While many diffusion filters rely on nonlinear concepts and may even require singular diffusivities or anisotropic diffusion tensors, our osmosis models show their high versatility already within a purely linear setting. Moreover, unlike gradient domain methods, osmosis is intrinsically invariant under global multiplicative changes of the grey values. In view of their promising potential, it is our hope that osmosis modelling will become a widely applied framework for visual computing.

Obviously our paper can only serve as a starting point, and there are many ways to continue research on osmosis processes for image analysis and synthesis. While our applications have exploited the nontrivial steady states of osmosis processes, it would also be interesting to see if the evolution itself has useful applications. On a theoretical side, we are working on a complete well-posedness and convergence theory for continuous linear osmosis processes, and we are also establishing a linear osmosis theory for semidiscrete and fully discrete processes; see [15] for first results. Last but not least, we will also consider nonlinear generalisations of our linear osmosis framework.

Acknowledgments. This work has been performed while all authors were members of the Mathematical Image Analysis Group. Financial support by the *Deutsche Forschungsgemeinschaft (DFG)* is gratefully acknowledged.

References

1. Borg, F.: What is osmosis? Explanation and understanding of a physical phenomenon. Technical report, Chydenius Institute, Jyväskylä University, Karleby, Finland (2003) arXiv:physics/0305011v1.
2. Weickert, J.: Anisotropic Diffusion in Image Processing. Teubner, Stuttgart (1998)
3. Hagenburg, K., Breuß, M., Vogel, O., Weickert, J., Welk, M.: A lattice Boltzmann model for rotationally invariant dithering. In Bebis, G., Boyle, R., Parvin, B., Koracin, D., Kuno, Y., Wang, J., Pajarola, R., Lindstrom, P., Hinkenjann, A., ao, M.L.E., Silva, C.T., Coming, D., eds.: Advances in Visual Computing. Volume 5876 of Lecture Notes in Computer Science. Springer, Berlin (2009) 949–959
4. Hagenburg, K., Breuß, M., Weickert, J., Vogel, O.: Novel schemes for hyperbolic PDEs using osmosis filters from visual computing. In Bruckstein, A.M., ter Haar Romeny, B., Bronstein, A.M., Bronstein, M.M., eds.: Scale Space and Variational Methods in Computer Vision. Volume 6667 of Lecture Notes in Computer Science. Springer, Berlin (2012) 532–543
5. Illner, R., Neunzert, H.: Relative entropy maximization and directed diffusion equations. *Mathematical Methods in the Applied Sciences* **16** (1993) 545–554
6. Frankot, R., Chellappa, R.: A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **10**(4) (July 1988) 439–451
7. Finlayson, G.D., Hordley, S.D., Drew, M.S.: Removing shadows from images. In Heyden, A., Sparr, G., Nielsen, M., Johansen, P., eds.: *Computer Vision – ECCV 2002*. Volume 2353 of Lecture Notes in Computer Science. Springer, Berlin (2002) 823–836
8. Morel, J.M., Petro, A.B., Sbert, C.: A PDE formalisation of retinex theory. *IEEE Transactions on Image Processing* **19**(11) (2010) 2825–2837
9. Fattal, R., Lischinski, D., Werman, M.: Gradient domain high dynamic range compression. In: *Proc. SIGGRAPH 2002*, San Antonio, TX (July 2002) 249–256
10. Pérez, P., Gagnat, M., Blake, A.: Poisson image editing. *ACM Transactions on Graphics* **22**(3) (2003) 313–318
11. Georgiev, T.: Covariant derivatives and vision. In Bischof, H., Leonardis, A., Pinz, A., eds.: *Computer Vision – ECCV 2006, Part IV*. Volume 3954 of Lecture Notes in Computer Science. Springer, Berlin (2006) 56–69
12. Risken, H.: *The Fokker–Planck Equation*. Springer, New York (1984)
13. Sochen, N.A.: Stochastic processes in vision: From Langevin to Beltrami. In: *Proc. Eighth International Conference on Computer Vision*. Volume 1., Vancouver, Canada, IEEE Computer Society Press (July 2001) 288–293
14. Wang, H., Hancock, E.R.: Probabilistic relaxation labelling using the Fokker–Planck equation. *Pattern Recognition* **41**(11) (November 2008) 3393–3411
15. Vogel, O., Hagenburg, K., Weickert, J., Setzer, S.: A fully discrete theory for linear osmosis filtering. In Kuijper, A., Bredies, K., Pock, T., Bischof, H., eds.: *Scale Space and Variational Methods in Computer Vision*. Volume 7893 of Lecture Notes in Computer Science. Springer, Berlin (2013) 368–379
16. Carlsson, S.: Sketch based coding of grey level images. *Signal Processing* **15** (1988) 57–83
17. Mainberger, M., Bruhn, A., Weickert, J., Forchhammer, S.: Edge-based compression of cartoon-like images with homogeneous diffusion. *Pattern Recognition* **44**(9) (2011) 1859–1873

18. Shor, Y., Lischinski, D.: The shadow meets the mask: Pyramid-based shadow removal. *Computer Graphics Forum* **22**(2) (April 2008) 577–586
19. Salamati, N., Germain, A., Süssstrunk, S.: Removing shadows from images using color and near-infrared. In: *Proc. 2011 IEEE International Conference on Image Processing*, Brussels, Belgium (September 2011) 1713–1716