

GEOMETRIC NUMERICAL INTEGRATION OF THE ASSIGNMENT FLOW

ALEXANDER ZEILMANN, FABRIZIO SAVARINO, STEFANIA PETRA, CHRISTOPH SCHNÖRR

ABSTRACT. The *assignment flow* is a smooth dynamical system that evolves on an elementary statistical manifold and performs contextual data labeling on a graph. We derive and introduce the *linear assignment flow* that evolves nonlinearly on the manifold, but is governed by a linear ODE on the tangent space. Various numerical schemes adapted to the mathematical structure of these two models are designed and studied, for the geometric numerical integration of both flows: embedded Runge-Kutta-Munthe-Kaas schemes for the nonlinear flow, adaptive Runge-Kutta schemes and exponential integrators for the linear flow. All algorithms are parameter free, except for setting a tolerance value that specifies adaptive step size selection by monitoring the local integration error, or fixing the dimension of the Krylov subspace approximation. These algorithms provide a basis for applying the assignment flow to machine learning scenarios beyond supervised labeling, including unsupervised labeling and learning from controlled assignment flows.

CONTENTS

1. Introduction	2
2. The Assignment Flow	4
3. Geometric Runge-Kutta Integration	6
3.1. General Approach	6
3.2. Application to the Assignment Flow	7
4. Linear Assignment Flow, Exponential Integrator	9
4.1. Linear Assignment Flow	9
4.2. Exponential Integrator	11
5. Step Sizes, Adaptivity	12
5.1. Linear Assignment Flow	13
5.2. Nonlinear Assignment Flow	16
6. Experiments and Discussion	17
6.1. Implementation Details	17
6.1.1. Ground Truth Flows	17
6.1.2. Termination criterion	18
6.1.3. Data	18
6.2. Nonlinear Flow: Embedded RKMK-Schemes	18
6.3. Linear Assignment Flow	20
6.3.1. Approximation Property	20
6.3.2. Adaptive RK Schemes	23
6.3.3. Exponential Integrator	24
7. Conclusion	25
References	26

Key words and phrases. image labeling, assignment flow, assignment manifold, geometric integration, adaptive step size selection, Krylov subspace approximation.

1. INTRODUCTION

Overview, motivation. The *assignment flow*, recently introduced by [ÅPSS17] and detailed in Section 2, denotes a *smooth dynamical system* evolving on an elementary statistical manifold, for the contextual classification of a finite set of data given on an arbitrary graph. Vertices of the graph are associated with the elements of the data set and correspond to locations in space and/or time in typical applications. *Classification* means to assign each datum exactly to one class representative, called *label*, out of a finite set of predetermined labels. *Contextual* classification means that these decision directly depend on each other, as encoded by the edges (adjacency relation) of the underlying graph. In the context of image analysis, classifying given image data on a regular grid graph in this way is called the *image labeling problem*. We point out, however, that the assignment flow applies to arbitrary data represented on a graph.

A key property of the assignment flow is that decision variables do *not* live in the space used to model the data. Rather, a probability simplex is associated with each datum, on which a flow evolves until it converges to one of the vertices of the simplex that encode the labels. Each simplex is equipped with the Fisher-Rao metric which turns the relative interior of the simplex into a smooth Riemannian manifold. It is this particular geometry that effectively promotes discrete decisions that interact in a *smooth* way. Replacing in addition the Riemannian (Levi Civita) connection by the α -connection (with $\alpha = 1$) introduced by Amari and Chentsov [AN00], enables to carry out basic geometric operations in a computationally *efficient* way. Keeping the assignment flow as ‘inference engine’ separate from the data space and model allows to flexibly apply it to a broad range of contextual data classification problems. We refer to [AN00, AJLS17] as basic texts on information geometry and to [KAH⁺15] for more information on the image labeling problem.

From a more distant viewpoint, our work ties in with the recent trend to explore the mathematics of deep networks from a dynamical systems perspective [E17]. A frequently cited paper in this respect is [HZRS16] where a connection was made between the so-called residual architecture of networks and explicit Euler integration steps of a corresponding system of nonlinear *ordinary differential equations (ODEs)*. We refer to [HR17] for a good exposition. While this offers a novel and fresh perspective on the *learning problem* of network parameters, it does not alter the basic ingredients of such networks that apparently have been adopted in an ad-hoc way, like parametrized static layers connected by nonlinear transition functions, ReLU activations etc.

By contrast, the assignment flow provides a smooth dynamical system on a graph (network), where all ingredients coherently fit into the overall mathematical framework. Based on this, we recently showed how discrete graphical models for image labeling can be evaluated using the assignment flow [HSÅS18], and how *unsupervised* labeling can be modeled by coupling the assignment flow and Riemannian gradient flows for label evolution on feature manifolds [ZZr⁺18]. Our current work, to be reported elsewhere, studies machine learning problems based on *controlling* the assignment flow. Here, in particular, *algorithms* play a decisive role that accurately *integrate the assignment flow numerically on the manifold* where it evolves. A thorough study of such algorithms is the subject of the present paper.

Contribution, organization. This paper presents two interrelated contributions, as illustrated by Figure 1.1.

- (1) We derive from the assignment flow – henceforth called *nonlinear assignment flow* – the *linear assignment flow*, that still is nonlinear but governed by a *linear ODE* on the tangent space. This property is attractive for modeling tasks (e.g. parameter estimation and control) as well as for the design of numerical algorithms. In particular, our experiments show that the linear flow closely approximates the nonlinear flow, as far as concerns the final *labeling* results.
- (2) We study a range of algorithms for numerically integrating both the nonlinear and the linear assignment flow, respectively, while properly taking into account the underlying *geometry*.
 - (a) Regarding the *nonlinear* assignment flow, we adopt the machinery of Lie group methods for the numerical integration of ODEs on manifolds [IMKNZ05] and devise corresponding extensions of classical Runge-Kutta schemes, called *RKMK schemes* (Runge-Kutta-Munthe-Kaas) in the literature [MK99]. We combine pairs of these extensions to form *embedded* RKMK schemes for adaptive step size control, analogous to classical embedded Runge-Kutta (RK) schemes [HNW08].
 - (b) Regarding the *linear* assignment flow, we take advantage in two alternative ways of the *linearity* of the flow on the tangent space.

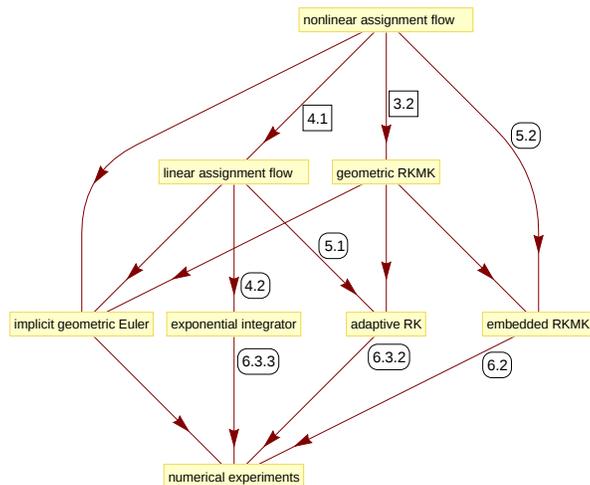


FIGURE 1.1. Topics addressed in this paper and their interrelations. Edge labels refer to the corresponding sections. Section numbers framed by squares address modeling aspects, whereas those framed by rounded squares address the design of algorithms and their numerical evaluation. Unlabelled edges mean ‘is derived from’ or ‘provides the basis for’.

- (i) On the one hand, we derive a local error estimate in order to apply classical Runge-Kutta schemes [HNW08] on the tangent space, with step sizes that adapt automatically.
- (ii) On the other hand, we evaluate the integral representation of the linear flow, due to Duhamels formula, and approximately evaluate this integral using Krylov subspace methods, as has been developed in the literature on exponential integrators [Saa92, HL97, HO10].

All these *explicit* numerical schemes are evaluated and discussed in Section 6, using ‘ground truth’ flows as a baseline that were computed using the *implicit* geometric Euler scheme with a sufficiently small step size. All algorithms are parameter free, except for specifying a single tolerance value with respect to the local error, that governs adaptive step size selection. Our experiments indicate a value for this parameter that ‘works’ regarding integration accuracy and labeling quality, but is not too conservative (i.e. small). In the case of the exponential integrator, we merely have to supply the final point of time T at which the linear assignment flow should be evaluated, in addition to the dimension of the Krylov subspace which controls the quality of the approximation. We conclude with a synopsis of our results in Section 7.

Basic notation. Index sets I and J index vertices $i \in I$ of the underlying graph and labels $j \in J$, respectively. \mathcal{S} and \mathcal{W} denote the basic statistical manifolds that we work with, defined in Section 2. Points $p, q \in \mathcal{S}$ are strictly positive probability vectors, and we denote efficiently by

$$qp = (q_1 \cdot p_1, \dots, q_{|J|} \cdot p_{|J|})^\top, \quad \frac{q}{p} = \left(\frac{q_1}{p_1}, \dots, \frac{q_{|J|}}{p_{|J|}} \right)$$

componentwise multiplication for general vectors, and componentwise subdivision only if $p \in \mathcal{S}$. Likewise, functions like the exponential function and the logarithm with vectors as arguments apply componentwise,

$$e^v = (e^{v_1}, e^{v_2}, \dots)^\top, \quad \log v = (\log v_1, \log v_2, \dots)^\top.$$

Exp and exp denote exponential mappings defined in Section 2, whereas expm denotes the *matrix* exponential in Section 4.2. The ordinary exponential function defined on the real line \mathbb{R} is always denoted by e^x , $x \in \mathbb{R}$.

$\mathbb{1} = (1, \dots, 1)^\top$ denotes the constant 1-vector with appropriate number of components depending on the context. We use the common shorthand $[n] = \{1, 2, \dots, n\}$ with $n \in \mathbb{N}$. $\|\cdot\|$ denotes the ℓ_2 -norm and $\|\cdot\|_p$ the ℓ_p -norm if $p \neq 2$.

2. THE ASSIGNMENT FLOW

We summarize the assignment flow introduced by [ÅPSS17] and related concepts required in this paper. Let $G = (I, E)$ be a given undirected graph and let

$$\mathcal{F}_I = \{f_i : i \in I\} \subset \mathcal{F} \quad (2.1a)$$

be data given in a metric space

$$(\mathcal{F}, d). \quad (2.1b)$$

We call \mathcal{F}_I *image data* even though the f_i typically represent *features* extracted from raw image at pixel $i \in I$ as a preprocessing step. G may be a regular grid graph as in low-level image processing or a less structured graph, with arbitrary connectivity in terms of the neighborhoods

$$\mathcal{N}_i = \{k \in I : ik = ki \in E\}. \quad (2.2)$$

We associate with each neighborhood \mathcal{N}_i weights satisfying

$$w_{ik} > 0, \quad \sum_{k \in \mathcal{N}_i} w_{ik} = 1, \quad \forall i \in I. \quad (2.3)$$

These weights parametrize the regularization property of the assignment flow and are assumed to be given. How to learn them from data in order to control the assignment flow will be reported elsewhere.

Along with \mathcal{F}_I we assume prototypical data

$$\mathcal{G}_J = \{g_j \in \mathcal{F} : j \in J\} \quad (2.4)$$

to be given, henceforth called *labels*. Each label g_j represents the data of class j . *Image labeling* denotes the problem to assign class labels to image data depending on the local context encoded by the graph G . We refer to [HSÅS18] for more details and background on the image labeling problem.

Assignment of labels to data are represented by discrete probability distributions

$$W_i = (W_{i1}, \dots, W_{iJ})^\top \in \mathcal{S}, \quad i \in I, \quad (2.5)$$

where

$$\mathcal{S} = \{p \in \mathbb{R}^{|J|} : p_j > 0, j \in J, \langle \mathbb{1}, p \rangle = 1\} \quad (2.6)$$

denotes the relatively open probability simplex equipped with the Fisher-Rao metric

$$g_p(u, v) = \sum_{j \in J} \frac{u_j v_j}{p_j}, \quad u, v \in T_0 = \{p \in \mathbb{R}^{|J|} : \langle \mathbb{1}, p \rangle = 0\}, \quad p \in \mathcal{S}, \quad (2.7)$$

which turns \mathcal{S} into a Riemannian manifold, with tangent spaces $T_p \mathcal{S} = T_0$ that do not depend on the base point p . In connection with \mathcal{S} , we define the

$$\mathbb{1}_{\mathcal{S}} = \frac{1}{|J|} (1, \dots, 1)^\top \in \mathbb{R}^{|J|} \quad (\mathbf{barycenter}) \quad (2.8)$$

of \mathcal{S} , i.e. the uniform distribution, the orthogonal projection

$$\Pi_{T_0} : \mathbb{R}^{|J|} \rightarrow T_0, \quad \Pi_{T_0}(z) = (\text{Diag}(\mathbb{1}) - \mathbb{1}\mathbb{1}_{\mathcal{S}}^\top)z, \quad (2.9a)$$

and the linear mapping

$$\Pi_p : \mathbb{R}^{|J|} \rightarrow T_0, \quad \Pi_p(z) = (\text{Diag}(p) - pp^\top)z, \quad p \in \mathcal{S} \quad (2.9b)$$

satisfying

$$\Pi_p = \Pi_p \Pi_{T_0} = \Pi_{T_0} \Pi_p. \quad (2.10)$$

Adopting the α -connection with $\alpha = 1$ from information geometry as introduced by Amari and Chentsov [AN00, Section 2.3], [AJLS17], the exponential map based on the corresponding affine geodesics reads

$$\text{Exp}: \mathcal{S} \times T_0 \rightarrow \mathcal{S}, \quad (p, v) \mapsto \text{Exp}_p(v) = \frac{e^{\frac{v}{p}}}{\langle p, e^{\frac{v}{p}} \rangle} p \quad (2.11a)$$

with inverse [ÅPSS17, Appendix]

$$\text{Exp}^{-1}: \mathcal{S} \times \mathcal{S} \rightarrow T_0, \quad (p, q) \mapsto \text{Exp}_p^{-1}(q) = \Pi_p \log \frac{q}{p}. \quad (2.11b)$$

Specifically, we define

$$\text{exp}_p = \text{Exp}_p \circ \Pi_p: \mathbb{R}^{|J|} = T_0 \oplus \mathbb{R}\mathbb{1} \rightarrow \mathcal{S}, \quad z \mapsto \frac{pe^z}{\langle p, e^z \rangle}, \quad \forall p \in \mathcal{S} \quad (2.12a)$$

with inverse [ÅPSS17, Appendix]

$$\text{exp}_p^{-1}: \mathcal{S} \rightarrow T_0, \quad q \mapsto \Pi_{T_0} \log \frac{q}{p}. \quad (2.12b)$$

Remark 2.1. Calling (2.12b) the “inverse” map is justified by the fact that exp_p does not depend on any constant component $\mathbb{R}\mathbb{1} \in T_0^\perp$ of the argument vector z . Yet, we choose $\mathbb{R}^{|J|}$ as domain because exp_p will be applied to arbitrary distance vectors $D_i \in \mathbb{R}^{|J|}$ (cf. (2.17)) arising from given data, and the notation indicates that the implementation does not need to remove this component explicitly [ÅPSS17, Remark 4].

These mappings naturally extend to the collections of assignment vectors (2.5), regarded as points on the

$$\mathcal{W} = \mathcal{S} \times \cdots \times \mathcal{S} \quad (|I| \text{ times}) \quad (\text{assignment manifold}) \quad (2.13)$$

with tangent space

$$\mathcal{T}_0 = T_0 \times \cdots \times T_0 \quad (|I| \text{ times}) \quad (2.14)$$

and the corresponding mappings

$$\mathbb{1}_{\mathcal{W}} = (\mathbb{1}_{\mathcal{S}}, \dots, \mathbb{1}_{\mathcal{S}}) \in \mathcal{W} \quad (\text{barycenter}) \quad (2.15a)$$

$$\Pi_{T_0}(Z) = (\Pi_{T_0}(Z_1), \dots, \Pi_{T_0}(Z_{|I|})) \in \mathcal{T}_0, \quad W \in \mathcal{W}, \quad Z \in \mathbb{R}^{|I||J|} \quad (2.15b)$$

$$\Pi_W(Z) = (\Pi_{W_1}(Z_1), \dots, \Pi_{W_{|I|}}(Z_{|I|})) \in \mathcal{T}_0, \quad W \in \mathcal{W}, \quad Z \in \mathbb{R}^{|I||J|} \quad (2.15c)$$

$$\text{Exp}_W(V) = (\text{Exp}_{W_1}(V_1), \dots, \text{Exp}_{W_{|I|}}(V_{|I|})) \in \mathcal{W}, \quad W \in \mathcal{W}, \quad V \in \mathcal{T}_0 \quad (2.15d)$$

and $\text{exp}_W, \text{Exp}_W^{-1}, \text{exp}_W^{-1}$ similarly defined based on (2.11b), (2.12a) and (2.12b). Finally, we define the *geometric mean* of assignment vectors [ÅPSS17, Lemma 5]

$$\mathcal{G}_i^w(W) = \text{Exp}_{W_i} \left(\sum_{k \in \mathcal{N}_i} w_{ik} \text{Exp}_{W_i}^{-1}(W_k) \right) = \text{exp}_{W_i} \left(\log \frac{\prod_{k \in \mathcal{N}_i} W_k^{w_{ik}}}{W_i} \right), \quad i \in I. \quad (2.16)$$

Using this setting, the assignment flow accomplishes image labeling as follows. Based on (2.1), (2.4), distance vectors

$$D = (D_1, \dots, D_{|I|}) \in \mathbb{R}^{|I||J|} \quad (\text{distance vectors}) \quad (2.17a)$$

$$D_i = (d(f_i, g_1), \dots, d(f_i, g_{|J|}))^\top, \quad i \in I \quad (2.17b)$$

are defined and mapped to

$$L(W) = \text{exp}_W(D) \in \mathcal{W}, \quad (\text{likelihood vectors}) \quad (2.18a)$$

$$L_i(W_i) := \frac{W_i e^{-\frac{1}{\rho} D_i}}{\langle W_i, e^{-\frac{1}{\rho} D_i} \rangle}, \quad \rho > 0, \quad i \in I, \quad (2.18b)$$

where ρ is a user parameter to normalize the distances induced by the specific features f_i at hand. This representation of the data is regularized by local geometric smoothing to obtain

$$S(W) \in \mathcal{W}, \quad S_i(W) = \mathcal{G}_i^w(L(W)), \quad i \in I, \quad (\text{similarity vectors}) \quad (2.19)$$

which in turn evolves the assignment vectors W_i , $i \in I$ through the

$$\dot{W} = \Pi_W(S(W)), \quad W(0) = \mathbb{1}_W. \quad (\text{assignment flow}) \quad (2.20)$$

Methods for numerically integrating this flow are examined in the following sections.

3. GEOMETRIC RUNGE-KUTTA INTEGRATION

We apply the general approach of [MK99] to our problem. For background and more details, we refer to [IMKNZ05] and [HLW06, Chapter 4].

3.1. General Approach. In order to apply Lie group methods to the integration of an ODE on a manifold \mathcal{M} , one has to check first if the ODE can be represented properly. Let

$$\Lambda: \mathbb{G} \times \mathcal{M} \rightarrow \mathcal{M} \quad (3.1a)$$

denote the action of a Lie group \mathbb{G} on \mathcal{M} satisfying

$$\Lambda(e, p) = p \quad \text{with identity } e \in \mathbb{G}, \quad (3.1b)$$

$$\Lambda(g_1 \cdot g_2, p) = \Lambda(g_1, \Lambda(g_2, p)), \quad \text{for all } g_1, g_2 \in \mathbb{G}, p \in \mathcal{M}. \quad (3.1c)$$

Furthermore, let \mathfrak{g} denote the Lie algebra of \mathbb{G} , $\mathfrak{X}(\mathcal{M})$ the set of all smooth vector fields on \mathcal{M} ,

$$\lambda: \mathfrak{g} \times \mathcal{M} \rightarrow \mathfrak{X}(\mathcal{M}) \quad (3.2)$$

a smooth function and λ_* the induced map defined by

$$\lambda_*: \mathfrak{g} \rightarrow \mathfrak{X}(\mathcal{M}), \quad (\lambda_*v)_p = \left. \frac{d}{dt} \lambda(tv, p) \right|_{t=0} \quad \text{for all } v \in \mathfrak{g}, p \in \mathcal{M}. \quad (3.3)$$

Then λ is a Lie algebra action if the induced map λ_* is a Lie algebra homomorphism, i.e. λ_* is linear and satisfies $\lambda_*[u, v] = [\lambda_*u, \lambda_*v]$, $u, v \in \mathfrak{g}$, with the Lie brackets on \mathfrak{g} and $\mathfrak{X}(\mathcal{M})$ on the left-hand side and the right-hand side, respectively. In particular, based on a Lie group action Λ , a Lie algebra action is given by [MK99, Lemma 4]

$$\lambda(v, p) = \Lambda(\exp_{\mathbb{G}}(v), p), \quad (3.4)$$

where $\exp_{\mathbb{G}}: \mathfrak{g} \rightarrow \mathbb{G}$ denotes the exponential map of \mathbb{G} . Thus, for this choice of λ , the induced map (3.3) is given by [MK99, Thm. 5]

$$(\lambda_*v)_p = \left. \frac{d}{dt} \Lambda(\exp_{\mathbb{G}}(tv), p) \right|_{t=0} \quad \text{for all } v \in \mathfrak{g}, p \in \mathcal{M}. \quad (3.5)$$

Now, given an ODE on \mathcal{M} , the *basic assumption* underlying the application of Lie group methods is the existence of a function $f: \mathbb{R} \times \mathcal{M} \rightarrow \mathfrak{g}$ such that the ODE admits the representation

$$\dot{y} = (\lambda_*f(t, y))_y, \quad y(0) = p. \quad (3.6)$$

For sufficiently small t , the solution of (3.6) then can be parametrized as

$$y(t) = \lambda(v(t), p), \quad (3.7a)$$

where $v(t) \in \mathfrak{g}$ satisfies the ODE

$$\dot{v} = (\text{dexp}_{\mathbb{G}}^{-1})_v(f(t, \lambda(v, p))), \quad v(0) = 0, \quad (3.7b)$$

with the inverse differential $(\text{dexp}_{\mathbb{G}}^{-1})_v$ of $\exp_{\mathbb{G}}$ evaluated at $v \in \mathfrak{g}$. A major advantage of the representation (3.7) is that the task of numerical integration concerns the ODE (3.7b) evolving on the vector space \mathfrak{g} , rather than the original ODE evolving on the manifold \mathcal{M} . As a consequence, established methods can be applied to (3.7b), possibly after approximating $\text{dexp}_{\mathbb{G}}^{-1}$ by a truncated series in a computationally feasible form.

3.2. Application to the Assignment Flow. Assume an ODE on \mathcal{S} defined by (2.6) is given. The application of the approach of Section 3.1 is considerably simplified by identifying $G = T_0$ with the *flat* tangent space (2.7) and consequently also $T_0 \cong \mathfrak{g} = T_e G$. One easily verifies that the action $\Lambda: T_0 \times \mathcal{S} \rightarrow \mathcal{S}$ defined as

$$\Lambda(v, p) = \exp_p(v), \quad (3.8)$$

with the right-hand side given by (2.12a), satisfies (3.1), i.e.

$$\Lambda(0, p) = p, \quad (3.9a)$$

$$\Lambda(v_1 + v_2, p) = \Lambda(v_1, \Lambda(v_2, p)) = \frac{pe^{v_1+v_2}}{\langle p, e^{v_1+v_2} \rangle}. \quad (3.9b)$$

Proposition 3.1. *The solution $W(t)$ to assignment flow (2.20) emanating from any $W_0 = W(0)$ admits the representation*

$$W(t) = \exp_{W_0}(V(t)) \quad (3.10a)$$

where $V(t) \in \mathcal{T}_0$ solves

$$\dot{V} = \Pi_{\mathcal{T}_0} S(\exp_{W_0}(V)), \quad V(0) = 0. \quad (3.10b)$$

Proof. Since geodesics through $0 \in T_0$ in directions $v \in T_0$ have the form $\gamma(t) = tv$, the differential of the exponential map of $T_0 = G$, $\exp_{T_0}(v) = \gamma(1) = v$, is the identity and thus (3.5) gives

$$(\lambda_* v)_p = \frac{d}{dt} \Lambda(\gamma(t), p) \Big|_{t=0} = \Pi_p(v), \quad (3.11)$$

with Π_p defined by (2.9b). As a result, the basic assumption (3.6) concerns ODEs on \mathcal{S} that admit the representation

$$\dot{p} = \Pi_p(f(t, p)), \quad p(0) = p_0, \quad (3.12)$$

for some function $f: \mathbb{R} \times \mathcal{S} \rightarrow T_0$ and some $p_0 \in \mathcal{S}$. Since $\lambda = \Lambda$ by (3.4), the parametrization (3.7) reads

$$p(t) = \Lambda(v(t), p_0) \quad (3.13a)$$

where $v(t) \in T_0$ solves

$$\dot{v} = f(t, \Lambda(v, p_0)), \quad v(0) = 0. \quad (3.13b)$$

This setting extends to the assignment flow by defining (cf. (2.15)) $\Lambda: \mathcal{T}_0 \times \mathcal{W} \rightarrow \mathcal{W}$ and $\lambda_*: \mathcal{T}_0 \rightarrow \mathcal{T}_0$ as

$$\Lambda(V, W) = \exp_W(V), \quad (\lambda_*(V))_W = \Pi_W(V). \quad (3.14)$$

The basic assumption (3.6) then reads

$$\dot{W} = (\lambda_* f(t, W))_W = \Pi_W(\Pi_{\mathcal{T}_0} S(W)) \stackrel{(2.10)}{=} \Pi_W(S(W)), \quad W(0) = \mathbb{1}_W, \quad (3.15)$$

which is the assignment flow (2.20). Due to (3.6), for any $W_0 = W(0)$, it admits the representation

$$W(t) = \Lambda(V(t), W_0), \quad (3.16a)$$

where $V(t) \in \mathcal{T}_0$ solves

$$\dot{V} = f(t, \Lambda(V, W_0)) = \Pi_{\mathcal{T}_0} S(\exp_{W_0}(V)), \quad V(0) = 0, \quad (3.16b)$$

which is (3.10). \square

Remark 3.2. While the basic formulation (2.20) of the assignment flow is *autonomous*, we keep in what follows the explicit time dependency of the function $f(t, \cdot)$ of the parametrization (3.10), because in more advanced scenarios the flow may become *non-autonomous*. A basic example concerns *unsupervised* problems [ZZr⁺18] where labels *vary*, and hence the distance vectors (2.17) and in turn the vector field defining the assignment flow depend on t .

Using the above representation and taking into account the simplifications of the general approach of Section 3.1, the **RKMK algorithm** [MK99] for integrating the assignment flow from $t = 0$ to $t = h$ is specified as follows. Let $a_{i,j}, b_j$ be the coefficients of an s -stage, q -th order classical Runge-Kutta method satisfying the consistency condition $c_i = \sum_{j \in [s]} a_{i,j}$ [HNW08, Section II]. Starting from any point

$$W_0 = W(0), \quad (3.17a)$$

the algorithm amounts to compute the vector fields

$$U^i = h \sum_{j \in [s]} a_{i,j} \tilde{U}^j, \quad i \in [s] \quad (3.17b)$$

$$\tilde{U}^i = f(hc_i, \Lambda(U^i, W_0)), \quad i \in [s] \quad (3.17c)$$

$$V = h \sum_{j \in [s]} b_j \tilde{U}^j \quad (3.17d)$$

and the update

$$W(h) = \Lambda(V, W_0). \quad (3.17e)$$

Replacing $W_0 \leftarrow W(h)$, computing the update and iterating this procedure generates the sequence $(W^{(k)})_{k \geq 0}$ which approximates $(W(t_k))_{k \geq 0}$, $t_k = kh$.

A s -stage *RKMK scheme* is specified using the corresponding Butcher tableau of the form

$$\begin{array}{c|cccc} 0 & & & & \\ c_2 & a_{21} & & & \\ c_3 & a_{31} & a_{32} & & \\ \vdots & \vdots & \vdots & \ddots & \\ c_s & a_{s1} & a_{s2} & \cdots & a_{s(s-1)} \\ \hline & b_1 & b_2 & \cdots & b_{s-1} & b_s \end{array}$$

Specifically, we consider the following *explicit RKMK schemes* of order 1, 2, 3, 4 ('FE' stands for Forward Euler):

$$\begin{array}{ccc} \begin{array}{c|c} 0 & \\ \hline 1 & 1 \end{array} & \begin{array}{c|cc} 0 & & \\ 1 & 1 & \\ \hline & 1/2 & 1/2 \end{array} & \begin{array}{c|ccc} 0 & & & \\ 1/3 & 1/3 & & \\ 2/3 & 0 & 2/3 & \\ \hline & 1/4 & 0 & 3/4 \end{array} & \begin{array}{c|ccc} 0 & & & \\ 1/2 & 1/2 & & \\ 1/2 & 0 & 1/2 & \\ 1 & 0 & 0 & 1 \\ \hline & 1/6 & 1/3 & 1/3 & 1/6 \end{array} \\ \text{Euler (FE)} & \text{Heun-2 (H2)} & \text{Heun-3 (H3)} & \text{Classical RK (RK4)} \end{array}$$

Note the increasing number of stages that raise the approximation order. This comes at a price, however, because each stage evaluates at step (3.17c) the right-hand side of (3.10b) which is the most expensive operation. As a consequence, it is not clear a priori if using a multi-stage scheme and a larger step size h is superior to a simpler scheme that is evaluated more frequently using a smaller step size.

In addition to the above explicit schemes, we consider the simplest *implicit RKMK scheme* ('BE' stands for Backward Euler)

$$\begin{array}{c|c} 1 & 1 \\ \hline & 1 \end{array} \\ \text{Euler (BE)}$$

Implicit schemes are known to be stable for much larger step sizes. Yet, they require to solve at every step a fixed point equation which is done by an iterative inner loop.

The performances of these numerical schemes are examined in Section 6.

4. LINEAR ASSIGNMENT FLOW, EXPONENTIAL INTEGRATOR

The ODE (3.10b) which parametrizes the assignment flow together with (3.10a), evolves on a linear space but is a *nonlinear* system of ordinary differential equations. In this section, we provide an *approximate representation* of the assignment flow within time intervals through a *linear* ODE evolving on the tangent space (Section 4.1), and a corresponding numerical scheme (Section 4.2).

The resulting flow on the assignment *manifold* is still nonlinear, though. The basic idea is to capture locally a major part of the nonlinearity of the (full) assignment flow, by a linear ODE on the tangent space that enables to apply alternative integration schemes.

4.1. Linear Assignment Flow. Our Ansatz has two ingredients. Firstly, we adopt the parametrization

$$W(t) = \text{Exp}_{W_0}(V(t)), \quad V(t) \in \mathcal{T}_0 \quad (4.1)$$

of the solution $W(t)$ to the assignment flow by a trajectory in the tangent space \mathcal{T}_0 , similar to (3.10a), except for using the ‘true’ exponential map (2.11a) and (2.15d), respectively, corresponding to the underlying affine connection. Secondly, we use an *affine approximation* of the vector field on the right-hand side of (2.20), that defines the assignment flow. The following corresponding definition generalizes the flow studied by [AE05] from the barycenter to arbitrary base points W_0 , and from a flow on \mathcal{S} to a flow on \mathcal{W} .

Definition 4.1 (linear assignment flow). We call *linear assignment flow* every flow induced by an ODE of the form

$$\dot{W} = \Pi_W \left(s_0 + S_0 \Pi_{W_0} \log \frac{W}{W_0} \right), \quad W(0) = W_0 \in \mathcal{W}, \quad (4.2)$$

with a fixed vector s_0 and a fixed matrix S_0 , for arbitrary W_0 .

An important property of the flow (4.2) – which explains its name – is the possibility to parametrize it by a *linear* ODE evolving on the tangent space \mathcal{T}_0 .

Proposition 4.2. *The linear assignment flow (4.2) admits the representation*

$$W(t) = \text{Exp}_{W_0}(V(t)), \quad (4.3a)$$

where $V(t) \in \mathcal{T}_0$ solves

$$\dot{V} = \Pi_{W_0}(s_0 + S_0 V), \quad V(0) = 0. \quad (4.3b)$$

Proof. Parametrization (4.1) yields

$$V(t) = \text{Exp}_{W_0}^{-1}(W(t)) \stackrel{(2.11b)}{=} \Pi_{W_0} \log \frac{W(t)}{W_0} \quad (4.4)$$

and by differentiation

$$\dot{V}(t) = \Pi_{W_0} \left(\frac{\dot{W}(t)}{W(t)} \right). \quad (4.5a)$$

Solving (4.2) for $\frac{\dot{W}}{W}$ after inserting (2.15c), and substitution in the preceding equation gives

$$= \Pi_{W_0} \left(s_0 + S_0 \text{Exp}_{W_0}^{-1}(W(t)) - \langle W(t), s_0 + S_0 \text{Exp}_{W_0}^{-1}(W(t)) \rangle \mathbb{1} \right), \quad (4.5b)$$

and since $\Pi_{W_0} \mathbb{1} = 0$ by (2.9b)

$$= \Pi_{W_0} \left(s_0 + S_0 \text{Exp}_{W_0}^{-1}(W(t)) \right) \stackrel{(4.4)}{=} \Pi_{W_0}(s_0 + S_0 V(t)). \quad (4.5c)$$

The initial condition follows from $V(0) = \text{Exp}_{W_0}^{-1}(W_0) = 0$. \square

Remark 4.3. Note that, despite the linearity of (4.3b), the resulting flow (4.3a) solving (4.2) is *nonlinear*. Thus, one may hope to capture the major nonlinearity of the full assignment flow (2.20) by a linear ODE on the tangent space, at least locally in some time interval. Within this interval, the evaluation of (2.19) is *not* required, and the linearity of the tangent space ODE (4.3b) can be exploited for integration.

We conclude this section by computing the natural choice

$$s_0 = S(W_0), \quad S_0 = dS_{W_0} \quad (4.6)$$

of the parameters of the linear assignment flow (4.2) in explicit form, where s_0 is immediate due to (2.19), but the Jacobian $S_0 = dS_{W_0}$ of $S(W)$, evaluated at W_0 , is not.

Proposition 4.4. *Let $S(W) \in \mathbb{R}^{|I||J|}$ denote the global similarity vector obtained by stacking the local similarity vectors $S_1(W), \dots, S_{|I|}(W)$ of (2.19). Then, with*

$$s_0 = S(W_0), \quad s_{0i} = S_i(W_0), \quad s_{0i,j} = S_{ij}(W_0) = (S_i(W_0))_j, \quad i \in I, j \in J \quad (4.7a)$$

and the projection $\Pi_{s_0,i}$ defined by (2.9b), the Jacobian of $S(W)$ at $W_0 \in \mathcal{W}$ is given by

$$S_0 = dS_{W_0} = \begin{pmatrix} A_{11}(W_0) & \cdots & A_{1|I|}(W_0) \\ \vdots & \ddots & \vdots \\ A_{|I|1}(W_0) & \cdots & A_{|I||I|}(W_0) \end{pmatrix} \in \mathbb{R}^{|I||J| \times |I||J|}, \quad (4.7b)$$

where each $|J| \times |J|$ block matrix has the form

$$A_{ik}(W_0)(V_k) = \begin{cases} w_{ik} \Pi_{s_0,i} \left(\frac{V_k}{W_{0k}} \right), & \text{if } k \in \mathcal{N}_i \\ 0, & \text{if } k \notin \mathcal{N}_i \end{cases}, \quad W_0 \in \mathcal{W}, \quad V_k \in T_0, \quad i, k \in I \quad (4.7c)$$

and the non-zero entries if $k \in \mathcal{N}_i$ (using (4.7a))

$$A_{ik,jl}(W_0) = (A_{ik}(W_0))_{jl} = w_{ik} \begin{cases} (1 - s_{0i,j}) \frac{s_{0i,j}}{W_{0k,j}}, & \text{if } j = l \\ -s_{0i,j} \frac{s_{0i,l}}{W_{0k,l}}, & \text{if } j \neq l \end{cases}, \quad j, l \in J. \quad (4.7d)$$

The proof follows below after two preparatory Lemmata.

Lemma 4.5. *Let $p \in \mathcal{S}$. Then the differential of $\exp_p : T_0 \rightarrow \mathcal{S}$ at $u \in T_0$ applied to $v \in T_0$ is given by*

$$d \exp_p(u)(v) = \Pi_{\exp_p(u)}(v), \quad u, v \in T_0, \quad p \in \mathcal{S}. \quad (4.8)$$

Moreover, we have

$$\exp_p(v - \log p) = \exp_{\mathbb{1}_{\mathcal{S}}}(v), \quad v \in T_0, \quad p \in \mathcal{S}. \quad (4.9)$$

Proof. Let $\gamma(t)$ be a smooth curve in T_0 with $\gamma(0) = u$ and $\dot{\gamma}(0) = v$. Using (2.12a), we compute

$$\left. \frac{d}{dt} \frac{pe^{\gamma(t)}}{\langle p, e^{\gamma(t)} \rangle} \right|_{t=0} = \frac{\langle p, e^u \rangle pve^u - \langle p, ve^u \rangle pe^u}{\langle p, e^u \rangle^2} = \frac{pe^u}{\langle p, e^u \rangle} \left(v - \frac{\langle pe^u, v \rangle}{\langle p, e^u \rangle} \mathbb{1} \right) = \Pi_{\exp_p(u)}(v), \quad (4.10)$$

which is (4.8). As for (4.9), using the representation

$$p \stackrel{(2.12a)}{=} \exp_{\mathbb{1}_{\mathcal{S}}}(\log p) = \exp_{\mathbb{1}_{\mathcal{S}}}(\Pi_{T_0} \log p), \quad (4.11)$$

where the last equation takes into account Remark (2.1), we obtain

$$\exp_p(v - \log p) = \exp_p(v - \Pi_{T_0} \log p) \stackrel{(4.11)}{=} \exp_{\exp_{\mathbb{1}_{\mathcal{S}}}(\Pi_{T_0} \log p)}(v - \Pi_{T_0} \log p) \quad (4.12a)$$

$$\stackrel{(3.8)}{=} \Lambda(v - \Pi_{T_0} \log p, \Lambda(\Pi_{T_0} \log p, \mathbb{1}_{\mathcal{S}})) \stackrel{(3.9)}{=} \Lambda(v - \Pi_{T_0} \log p + \Pi_{T_0} \log p, \mathbb{1}_{\mathcal{S}}) \quad (4.12b)$$

$$= \Lambda(v, \mathbb{1}_{\mathcal{S}}) = \exp_{\mathbb{1}_{\mathcal{S}}}(v). \quad (4.12c)$$

□

We use this Lemma to represent the similarity vectors in a convenient form for subsequently proving Prop. 4.4.

Lemma 4.6. *The similarity vectors (2.19) admit the representation*

$$S_i(W) = \exp_{\mathbb{1}_{\mathcal{S}}} \left(\sum_{k \in \mathcal{N}_i} w_{ik} \left(\log W_k - \frac{1}{\rho} D_k \right) \right), \quad i \in I. \quad (4.13)$$

Proof. By (2.19) and (2.16), we obtain

$$S_i(W) = \exp_{W_i} \left(\log \frac{\prod_{k \in \mathcal{N}_i} L_k(W_k)^{w_{ik}}}{W_i} \right) = \exp_{W_i} \left(\sum_{k \in \mathcal{N}_i} w_{ik} \log L_k(W_k) - \log W_i \right) \quad (4.14a)$$

$$\stackrel{(2.18b)}{=} \exp_{W_i} \left(\sum_{k \in \mathcal{N}_i} w_{ik} \left(\log W_k - \frac{1}{\rho} D_k - \log (\langle W_k, e^{-\frac{1}{\rho} D_k} \rangle \mathbb{1}) \right) - \log W_i \right) \quad (4.14b)$$

using again $\exp_p(v + \lambda \mathbb{1}) = \exp_p(v)$ for all $\lambda \in \mathbb{R}$, $v \in T_0$, $p \in \mathcal{S}$ (cf. Remark 2.1)

$$= \exp_{W_i} \left(\sum_{k \in \mathcal{N}_i} w_{ik} \left(\log W_k - \frac{1}{\rho} D_k \right) - \log W_i \right) \stackrel{(4.9)}{=} \exp_{\mathbb{1}_{\mathcal{S}}} \left(\sum_{k \in \mathcal{N}_i} w_{ik} \left(\log W_k - \frac{1}{\rho} D_k \right) \right). \quad (4.14c)$$

□

Proof of Prop. 4.4. Setting

$$S_i(W) \stackrel{(4.13)}{=} \exp_{\mathbb{1}_{\mathcal{S}}} \circ Z_i(W), \quad Z_i(W) = \sum_{k \in \mathcal{N}_i} w_{ik} \left(\log W_k - \frac{1}{\rho} D_k \right), \quad (4.15)$$

we compute using a smooth curve $\gamma(t)$ in \mathcal{W} with $\gamma(0) = W$ and $\dot{\gamma}(0) = V$,

$$dZ_i(W)(V) = \frac{d}{dt} Z_i(\gamma(t)) \Big|_{t=0} = \sum_{k \in \mathcal{N}_i} w_{ik} \frac{d}{dt} \log(\gamma(t)) \Big|_{t=0} = \sum_{k \in \mathcal{N}_i} w_{ik} \frac{V_k}{W_k}. \quad (4.16)$$

Thus, using (4.15) and (4.8) gives

$$dS_i(W)(V) \stackrel{(4.15)}{=} d \exp_{\mathbb{1}_{\mathcal{S}}} (Z_i(W)) (dZ_i(W)(V)) \stackrel{(4.8), (4.15)}{=} \Pi_{S_i(W)} (dZ_i(W)(V)), \quad (4.17a)$$

and using the linearity of the map $\Pi_{S_i(W)}$ and (4.16),

$$= \sum_{k \in \mathcal{N}_i} w_{ik} \Pi_{S_i(W)} \left(\frac{V_k}{W_k} \right), \quad (4.17b)$$

which proves (4.7c). Inserting $\Pi_{S_i(W)}$ due to (2.9b) yields (4.7d). □

The following section specifies an alternative integration scheme for the linear assignment flow (4.2). Its approximation properties are numerically examined in Section 6.

4.2. Exponential Integrator. We focus on the linear ODE (4.3b) that together with (4.3a) determines the linear assignment flow due to (4.2). The solution to (4.3b) is given by Duhamel's formula [Tes12],

$$V(t) = \expm(tA) V(0) + \int_0^t \expm((t-\tau)A) a \, d\tau \quad \text{where } A = \Pi_{W_0} S_0, \quad a = \Pi_{W_0} s_0, \quad (4.18)$$

which involves the matrix exponential of the matrix A of dimension $|I||J| \times |I||J|$ (square of number of pixels \times number of labels), which can be quite large in image labeling problems (10^4 – 10^7 variables). Explicitly computing the matrix exponential is neither feasible, because it is dense even if A is sparse, nor required in view of the multiplication with the vector a . Rather, taking into account $V(0) = 0$ and that uniformly converging series can be integrated term by term, we set $t = T$ large enough and evaluate

$$V(T) = \int_0^T \expm((T-\tau)A) a \, d\tau = \expm(TA) \int_0^T \sum_{k=0}^{\infty} \frac{(-\tau A)^k}{k!} a \, d\tau \quad (4.19a)$$

$$= \expm(TA) \sum_{k=0}^{\infty} \left[\frac{\tau (-\tau A)^k}{(k+1)!} \right]_{\tau=0}^T a = T \expm(TA) \sum_{k=0}^{\infty} \frac{(-TA)^k}{(k+1)!} a \quad (4.19b)$$

$$= T \varphi_1(TA) a \quad (4.19c)$$

where φ_1 is the entire function

$$\varphi_1(z) = \sum_{k=0}^{\infty} \frac{z^k}{(k+1)!} = \frac{e^z - 1}{z}, \quad (4.20)$$

whose series representation yields valid expressions (4.19c) also if the matrix A is singular.

We refer to [Hig08] for a detailed exposition of matrix functions and to [MVL03] and [Hig08, Sections 10 and 13] for a survey of methods for computing the matrix exponential and the product of matrix functions times a vector. For large problem sizes, the established methods of the two latter references are known to deteriorate, however, and methods based on Krylov subspaces have been developed [Saa92, HL97] and become the method of choice in connection with exponential integrators [HO10].

We confine ourselves with sketching below a state-of-the-art method [NW12] for the approximate numerical evaluation of (4.19). The evaluation of its performance for integrating the linear assignment flow and a comparison to the methods of Section 5.1, are reported in Section 6. A more comprehensive evaluation of further recent methods for evaluating (4.19) that cope with large problem sizes as well (e.g. [AMH11]), is beyond the scope of this paper.

In order to compute approximately $\varphi_1(TA)a$, one considers the Krylov subspace

$$K_m = \text{span}\{a, Aa, \dots, A^{m-1}a\}, \quad (4.21)$$

with orthogonal basis $V_m = (v_1, \dots, v_m)$ arranged as column vectors of an orthogonal matrix V_m and computed using the basic Arnoldi iteration [Saa92]. The action of A is approximated by

$$H_m = V_m^\top AV_m \quad (4.22)$$

which in turn yields the approximation

$$\varphi_1(A)a \approx \varphi_1(V_m H_m V_m^\top)a = V_m \varphi_1(H_m) V_m^\top a = \|a\| V_m \varphi_1(H_m) e_1, \quad (4.23)$$

where $e_1 = (1, 0, \dots, 0)^\top$ denotes the first unit vector and the last equality is implied by the Arnoldi iteration producing V_m, H_m , which sets $v_1 = a/\|a\|$. Note that φ_1 merely has to be applied to the much smaller $m \times m$ matrix H_m , which can be safely and efficiently computed using standard methods [MVL03, Hig08]. The vector of $\varphi_1(H_m)e_1$ can be recovered [Sid98, Thm. 1] in form of the upper m entries of the last column of $\text{expm}(\widehat{H}_{m,k})$ with the extended matrix

$$\widehat{H}_m = \begin{pmatrix} H_m & e_1 \\ 0 & 0 \end{pmatrix}. \quad (4.24)$$

If the degree of the minimal polynomial of a (i.e. the the nonzero monic polynomial p of lowest degree such that $p(A)a = 0$) is equal to m , then the approximation (4.23) is even exact [Saa92, Thm. 3.6].

5. STEP SIZES, ADAPTIVITY

We specify in this section step size selection for the numerical RKMK schemes of Section 3. In addition, for the *linear assignment flow* (Section 4.1), we conduct a local error analysis in Section 5.1 for RK schemes based on the linearity of the tangent space ODE that governs this flow. A corresponding explicit error estimate enables to determine a sequence $(h_k)_{k \geq 0}$ of step sizes that ensure a prespecified local accuracy at each step k .

In order to determine step sizes for the *nonlinear assignment flow* (2.20), we proceed differently, because the corresponding vector field depends nonlinearly on the current iterate and estimating local Lipschitz constants will be expensive and less sharp. We therefore adapt in Section 5.2 classical methods for local error estimation and step size selection for nonlinear ODEs based on *embedded* Runge-Kutta methods [HNW08, Section II.4], to the *geometric* RKMK methods of Section 3.

The experimental evaluation of both approaches is reported in Section 6.

5.1. Linear Assignment Flow. We focus on the linear ODE (4.3b) that together with (4.3a) determines the linear assignment flow (4.2). Due to its approximation property demonstrated in Section 6.3.1, we only consider the linearization point $W_0 = \mathbb{1}_{\mathcal{W}}$. Since the ODE (4.2) evolves on the linear space \mathcal{T}_0 , we apply the *classical* s -stage explicit Runge-Kutta (RK) scheme, rather than the *geometric* s -stage RKMK scheme (3.17), to obtain

$$U^i = \Pi_{W_0} s_0 + \Pi_{W_0} S_0 \left(V^{(k)} + h \sum_{j \in [s-1]} a_{i,j} U^j \right), \quad i \in [s], \quad (5.1a)$$

$$V^{(k+1)} = V^{(k)} + h \sum_{i \in [s]} b_i U^i, \quad V^{(0)} = V(0) = 0. \quad (5.1b)$$

Specifically, regarding the *explicit schemes* listed at the end of Section 3.2 in terms of their Butcher tableaus, consecutively inserting (5.1a) into (5.1b) yields with the shorthands a, A defined by (4.18),

$$V^{(k+1)} = ha + (I + hA)V^k, \quad (\mathbf{FE}) \quad (5.2a)$$

$$V^{(k+1)} = \left(h + \frac{h^2}{2}A \right) a + \left(I + hA + \frac{h^2}{2}A^2 \right) V^{(k)}, \quad (\mathbf{H2}) \quad (5.2b)$$

$$V^{(k+1)} = \left(h + \frac{h^2}{2}A + \frac{h^3}{6}A^2 \right) a + \left(I + hA + \frac{h^2}{2}A^2 + \frac{h^3}{6}A^3 \right) V^{(k)}, \quad (\mathbf{H3}) \quad (5.2c)$$

$$V^{(k+1)} = \left(h + \frac{h^2}{2}A + \frac{h^3}{6}A^2 + \frac{h^4}{24}A^3 \right) a + \left(I + hA + \frac{h^2}{2}A^2 + \frac{h^3}{6}A^3 + \frac{h^4}{24}A^4 \right) V^{(k)}, \quad (\mathbf{RK4}) \quad (5.2d)$$

with

$$V^{(0)} = 0. \quad (5.2e)$$

Comparison with (4.18) shows that due to the linearity of the ODE, each scheme results in a corresponding Taylor series approximation, depending on its order q , of the equation

$$V(t_{k+1}) = h\varphi_1(hA)a + \expm(hA)V(t_k) \quad (5.3)$$

that is,

$$V^{(k+1)} = h \left(\sum_{i=0}^{q-1} \frac{(hA)^i}{(i+1)!} \right) a + \left(\sum_{i=0}^q \frac{(hA)^i}{i!} \right) V^{(k)} \quad (5.4a)$$

$$= p_{1,q}(hA)a + p_{2,q}(hA)V^{(k)}, \quad (5.4b)$$

with matrix-valued polynomials $p_{1,q}, p_{2,q}$. Our strategy for choosing the step size h is based on the local error estimate specified below as Theorem 5.2 and prepared by the following Lemma.

Lemma 5.1. *Let $q \in \mathbb{N}$ and $t \in \mathbb{R}$. Then*

$$\sum_{i=0}^q \frac{t^i}{i!} = e^t \frac{\Gamma(1+q, t)}{q!} \quad (5.5)$$

with the incomplete Gamma function

$$\Gamma(1+q, t) = \int_t^\infty \tau^q e^{-\tau} d\tau. \quad (5.6)$$

Proof. Partial integration shows the recursion

$$\Gamma(1+q, t) = -\tau^q e^{-\tau} \Big|_t^\infty + \int_t^\infty q\tau^{q-1} e^{-\tau} d\tau \quad (5.7a)$$

$$= t^q e^{-t} + q\Gamma(q, t) \quad (5.7b)$$

$$= t^q e^{-t} + q(t^{q-1} e^{-t} + (q-1)\Gamma(q-1, t)) \quad (5.7c)$$

$$= (t^q + qt^{q-1})e^{-t} + q(q-1)\Gamma(q-1, t) \quad (5.7d)$$

$$= (t^q + qt^{q-1})e^{-t} + q(q-1)(t^{q-2} e^{-t} + (q-2)\Gamma(q-2, t)) \quad (5.7e)$$

$$= (t^q + qt^{q-1} + q(q-1)t^{q-2})e^{-t} + q(q-1)(q-2)\Gamma(q-2, t) \quad (5.7f)$$

$$= \dots \quad (5.7g)$$

$$= (t^q + qt^{q-1} + \dots + q(q-1) \cdots 2 \cdot t)e^{-t} + q! \underbrace{\Gamma(1, t)}_{=e^{-t}} \quad (5.7h)$$

$$= (t^q + qt^{q-1} + \dots + q(q-1) \cdots 2 \cdot t + q!)e^{-t}. \quad (5.7i)$$

Subdividing both sides by $q! e^{-t}$ yields (5.5). \square

Theorem 5.2. Let $V(t)$ solve (4.2) with $W_0 = \mathbb{1}_{\mathcal{W}}$, and let $(V^{(k)})_{k>0}$ be a sequence generated by a RK scheme (5.1) of order q . Set $V(t_k) = V^{(k)}$ in (5.3). Then $V(t_{k+1})$ in (5.3) is the exact value of the linear assignment flow emanating from $V^{(k)}$, and regarding (5.4) the local error estimate

$$\|V(t_{k+1}) - V^{(k+1)}\| \leq e^{h\|A\|} \left(1 - \frac{\Gamma(1+q, h\|A\|)}{q!}\right) \left(\frac{\|a\|}{\|A\|} + \|V^{(k)}\|\right) \quad (5.8a)$$

$$< e^{h\|A\|} (1 - e^{-h\|A\|})^{(1+q)} \left(\frac{\|a\|}{\|A\|} + \|V^{(k)}\|\right) \quad (5.8b)$$

holds, where $\Gamma(1+q, h\|A\|)$ is given by (5.6) and $\|A\|$ denotes the spectral norm of the matrix $A = \Pi_{W_0}(S_0)$.

Proof. Using (5.3), (5.4) and $V(t_k) = V^{(k)}$, we bound the local error by

$$\|V(t_{k+1}) - V^{(k+1)}\| \leq \|h\varphi_1(hA) - p_{1,q}(hA)\| \|a\| + \|\expm(hA) - p_{2,q}(hA)\| \|V^{(k)}\|, \quad (5.9a)$$

and inserting the series (5.4a) gives

$$\leq h \left(\sum_{i=q}^{\infty} \frac{(h\|A\|)^i}{(i+1)!} \right) \|a\| + \left(\sum_{i=q+1}^{\infty} \frac{(h\|A\|)^i}{i!} \right) \|V^{(k)}\|. \quad (5.9b)$$

Both series absolutely converge for any h . By Lemma 5.1, we have

$$\sum_{i=q}^{\infty} \frac{t^i}{(i+1)!} \stackrel{j=i+1}{=} \sum_{j=q+1}^{\infty} \frac{1}{t} \cdot \frac{t^j}{j!} = \frac{e^t}{t} \left(1 - \frac{\Gamma(1+q, t)}{q!}\right), \quad (5.10a)$$

$$\sum_{i=q+1}^{\infty} \frac{t^i}{i!} = e^t \left(1 - \frac{\Gamma(1+q, t)}{q!}\right). \quad (5.10b)$$

Applying these equations to (5.9b) yields

$$h \sum_{i=q}^{\infty} \frac{(h\|A\|)^i}{(i+1)!} = \frac{e^{h\|A\|}}{\|A\|} \left(1 - \frac{\Gamma(1+q, h\|A\|)}{q!}\right), \quad (5.11a)$$

$$\sum_{i=q+1}^{\infty} \frac{(h\|A\|)^i}{i!} = e^{h\|A\|} \left(1 - \frac{\Gamma(1+q, h\|A\|)}{q!}\right), \quad (5.11b)$$

and substitution into (5.9)

$$\|V(t_{k+1}) - V^{(k+1)}\| \leq e^{h\|A\|} \left(1 - \frac{\Gamma(1+q, h\|A\|)}{q!}\right) \left(\frac{\|a\|}{\|A\|} + \|V^{(k)}\|\right), \quad (5.12)$$

which is (5.8a). To show (5.8b), we use the representation

$$\frac{1}{p}\Gamma\left(\frac{1}{p}, x^p\right) \stackrel{(5.15)}{=} \frac{1}{p} \int_{x^p}^{\infty} t^{\frac{1}{p}-1} e^{-t} dt \stackrel{t=\tau^p}{=} \int_x^{\infty} e^{-\tau^p} d\tau \quad (5.13)$$

and the lower bound [Alz97, Corollary of Thm. 1]

$$\frac{1}{\Gamma(1+1/p)} \int_x^{\infty} e^{-t^p} dt > 1 - (1 - e^{-\alpha x^p})^{1/p}, \quad \alpha \geq \max\{1, (\Gamma(1+1/p))^{-p}\}, \quad (5.14)$$

that holds for all $x > 0$ and $0 < p \neq 1$, with the Gamma function

$$\Gamma(q) = \int_0^{\infty} \tau^{q-1} e^{-\tau} d\tau, \quad \Gamma(n+1) = n! \quad \text{if } n \in \mathbb{N}. \quad (5.15)$$

Put

$$x = h\|A\| \quad \text{and} \quad p = \frac{1}{1+q}. \quad (5.16)$$

Since $q \geq 1$ and $(\Gamma(1+1/p))^{-p} = \Gamma(2+q)^{-\frac{1}{1+q}} = \left(\frac{1}{(q+1)!}\right)^{\frac{1}{1+q}} < 1$, we set $\alpha = 1$ in view of (5.14). Furthermore, we have

$$\Gamma(1+q, h\|A\|) \stackrel{(5.16)}{=} \Gamma(1/p, x) = \Gamma(1/p, (x^{1/p})^p) \stackrel{(5.13)}{=} p \int_{x^{1/p}}^{\infty} e^{-t^p} dt \quad (5.17a)$$

$$\stackrel{(5.14), \alpha=1}{>} p\Gamma(1+1/p)(1 - (1 - e^{-x})^{1/p}), \quad (5.17b)$$

and using $p\Gamma(1+1/p) \stackrel{(5.16)}{=} \frac{\Gamma(2+q)}{1+q} = \frac{(1+q)!}{1+q} = q!$, since q is integer,

$$= q!(1 - (1 - e^{-x})^{(1+q)}). \quad (5.17c)$$

Thus,

$$e^x \left(1 - \frac{\Gamma(1/p, x)}{q!}\right) < e^x (1 - e^{-x})^{(1+q)} \quad (5.18)$$

which after substituting (5.16) and in turn into (5.12), proves (5.8b). \square

Theorem 5.2 enables to control the local integration error by choosing the step size h , using the simple form of the bound (5.8), depending on the constants $\|a\|$, $\|A\|$ and the norm $\|V^{(k)}\|$ of the current iterate. Specifically, we choose

$$h = h_k \quad \text{such that} \quad \|V(t_{k+1}) - V^{(k+1)}\| \leq \tau \quad (5.19)$$

by (5.8), for some prespecified value τ . Inspecting the parametrization (4.3) shows that $\|V(t)\|$ grows – and hence the step sizes (5.19) decrease – until $W(t)$ is close enough to a vertex of \mathcal{W} (which represents a labeling) and satisfies a termination criterion that stops the chosen iterative RK scheme (5.1).

In order to check how large $\|V(t)\|$ then will be, assume

$$W_i = (\varepsilon, \dots, \varepsilon, 1 - (|J| - 1)\varepsilon, \varepsilon, \dots, \varepsilon) \in \mathbb{R}^{|J|} \quad \text{and} \quad \varepsilon \ll \frac{1}{|J| - 1} \leq 1, \quad (5.20)$$

that is $W_{ij} \approx 1$ and $W_{il} \approx 0$ if $l \neq j$. Then with $W_{0i} = \mathbb{1}_S$ and by (4.3), (2.12a)

$$V_i = \text{Exp}_{\mathbb{1}_S}^{-1}(W_i) = \Pi_{\mathbb{1}_S}(\log W_i - \log \mathbb{1}_S) = \frac{1}{|J|} \left(\log W_i - \frac{1}{|J|} \langle \mathbb{1}, \log W_i \rangle \mathbb{1} \right) \quad (5.21a)$$

$$\log W_i \approx (\log \varepsilon, \dots, \log \varepsilon, 0, \log \varepsilon, \dots, \log \varepsilon), \quad \frac{1}{|J|} \langle \mathbb{1}, \log W_i \rangle \approx \frac{|J| - 1}{|J|} \log \varepsilon \approx \log \varepsilon \quad (5.21b)$$

and hence

$$\|V_i\| \approx \frac{1}{|J|} \log \frac{1}{\varepsilon}, \quad \|V\| \approx \frac{|I|}{|J|} \log \frac{1}{\varepsilon}. \quad (5.21c)$$

Thus, as soon as the norm $\|V(t)\|$ has grown to the order $\log \frac{1}{\varepsilon}$, a termination criterion that checks if $W(t)$ is ε -close to some vertex of \mathcal{W} , will be satisfied.

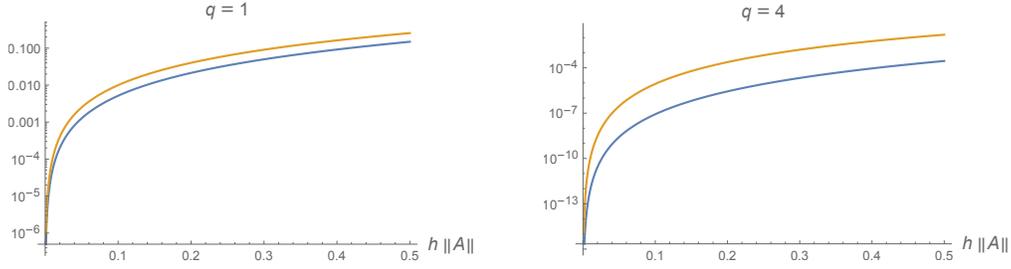


FIGURE 5.1. The factor $e^{h\|A\|}(1-e^{-h\|A\|})^{(1+q)}$ of the upper bound (5.8) (beige curve) and the numerically computed exact factor due to (5.12) (blue line), as a function of $h\|A\|$, for $q = 1$ (left panel) and $q = 4$ (right panel). The *relative* overestimation factor increases with the order q and leads to more or less conservative step size choices (5.19). Comparing the *absolute* ordinate values of both panels shows that in order to achieve (5.19), using a higher-order RK-scheme (5.1) enables to choose a larger step size.

Figure 5.1 quantitatively illustrates how much the factor $e^{h\|A\|}(1-e^{-h\|A\|})^{(1+q)}$ of the upper bound (5.8) overestimates the exact factor (5.12) computed in the proof of Thm. 5.2, and hence how conservative (i.e. too small) the step size h_k will be chosen to achieve (5.19). The curves of Figure 5.1 show that the estimate (5.8) is fairly tight and suited to adapt the step size. Furthermore, comparing the ordinate values of both panels for $q = 1$ and $q = 4$ shows that, in order to achieve a fixed accuracy τ in (5.19), using a *higher-order* RK scheme (5.1) enables to choose a *larger* step size.

5.2. Nonlinear Assignment Flow. Similar to the preceding section, we wish to select step sizes $(h_k)_{k \geq 0}$ in order to control the local error on the left-hand side of (5.8). Because an estimate like the right-hand side of (5.8) that is valid at each step k , is not available for the *nonlinear* assignment flow, we adapt established *embedded RK methods* [HNW08, Section II.V] to the geometric RKM schemes (3.17).

The basic strategy is to evaluate *twice* step (3.17d)

$$V = h \sum_{j \in [s]} b_j \tilde{U}^j, \quad \hat{V} = h \sum_{j \in [s]} \hat{b}_j \tilde{U}^j, \quad (5.22)$$

using a second collection of coefficients \hat{b}_j , $j \in s$, but with the *same* vector fields U^i, \tilde{U}^i , $i \in [s]$. Thus each embedded method can be specified by *augmenting* the corresponding Butcher tableau accordingly,

$$\begin{array}{c|cccc} 0 & & & & \\ c_2 & a_{21} & & & \\ c_3 & a_{31} & a_{32} & & \\ \vdots & \vdots & \vdots & \ddots & \\ c_s & a_{s1} & a_{s2} & \cdots & a_{s(s-1)} \\ \hline & b_1 & b_2 & \cdots & b_{s-1} & b_s \\ & \hat{b}_1 & \hat{b}_2 & \cdots & \hat{b}_{s-1} & \hat{b}_s \end{array}$$

Proper embedded methods combine a pair of RK schemes of *different* order q, \hat{q} so that $\|V - \hat{V}\|$ indicates if the step size h is small enough, at each step k of the overall iteration. Since the vectors U^i, \tilde{U}^i , $i \in [s]$ are used twice, this comes at little additional costs. We also point out that unlike the linear case, the magnitude $\|V\|$ of tangent vectors has much less influence, because the scheme (3.17) that is consecutively applied at each step k , is based on (3.10b) with the initial condition $V(0) = 0$. As a consequence, the magnitude $\|V\|$ of the update (3.17d) will be relatively small at each step k .

We list the tableaus of two embedded methods that we evaluate in Section 6.

0	
1	1
	1 0
	1/2 1/2

(RK-1/2)

0		
1/3	1/3	
2/3	0	2/3
	1/4	0 3/4
	1/3	2/3 0

(RK-3/2)

The first method combines the forward Euler scheme and Heun’s method of order 2. The second method complements Heun’s method of order 3 as specified on [HNW08, p. 166]. We call **RKMK-1/2** and **RKMK-3/2** the geometric versions of these schemes when they are applied in connection with (3.17).

We conclude this section by specifying the extension of (3.17) in order to include adaptive step size control. Fix parameters $0 < \tau \ll 1$, $n_\tau \in \mathbb{N}$ and set an initial sufficiently small step size $h = h_0$. At each step k , using the distance d_I defined by (6.1):

1. Compute $U^i, \tilde{U}^i, i \in [s]$.
 2. Compute V, \hat{V} by (5.22).
 3. If $d_I(V, \hat{V}) < \frac{\tau}{n_\tau}$, then increase the step size: $h \leftarrow 1.25h$, compute the update $W(h)$ by (3.17e), set $W_0 \leftarrow W(h)$ and proceed with the next iteration $k + 1$ and step 1. Otherwise continue with step 4.
 4. If $d_I(V, \hat{V}) < \tau$, then keep the step size h , compute the update $W(h)$ by (3.17e), set $W_0 \leftarrow W(h)$ and proceed with the next iteration $k + 1$ and step 1. Otherwise continue with step 5.
 5. Decrease the step size $h \leftarrow \frac{h}{2}$ and repeat iteration k , i.e. continue with step 1.
- (5.23)

Typical parameter values are $\tau = 0.01, n_\tau = 20$. Starting with an small initial step size h_0 , the algorithm adaptively generates a sequence (h_k) whose values increase whenever the local error estimate is much smaller (by a factor n_τ) than the prescribed tolerance τ .

6. EXPERIMENTS AND DISCUSSION

This section is organized as follows (see also Figure 1.1). We specify details of our implementation in Section 6.1. Section 6.2 reports the evaluation of the geometric RKMK schemes (3.17) with embedded step size control (5.23), for integrating the *nonlinear* assignment flow. Section 6.3 is devoted to the *linear* assignment flow: Assessment of how closely it approximates the nonlinear assignment flow, evaluation of the RK schemes (5.1) with adaptive step size selection (5.19), and evaluation of the exponential integrator introduced in Section 4.2.

6.1. Implementation Details. All algorithms were implemented and evaluated using Mathematica. In particular, we did *not* apply any assignment normalization as suggested by [ÅPSS17, Section 3.3.1], since Mathematica can work with arbitrary numerical precision. As a consequence, our results refute the claim of the authors of [BFPS17] that properties of the assignment flow are largely caused by this normalization step and a particular numerical scheme. Rather, our experiments illustrate *intrinsic* properties of the assignment flow as well as the reliability and efficiency of a *variety* of algorithms for integrating this flow numerically, conforming to the underlying geometry, as derived in the present paper.

Throughout the experiments, we used *uniform* weights in (2.19) and (2.16), respectively, since how to choose these ‘control variables’ in a proper way depending on the application at hand, is subject of our current work and will be reported elsewhere. Yet, we point out that the algorithms of the present paper *do* cover such more general scenarios. For example, the simplest geometric RKMK scheme (3.17) was recently used for integrating the assignment flow in *unsupervised* scenarios [ZZr+18], where labels evolve and hence distance vectors (2.17) no longer are static but vary with time $D = D(t)$, too.

6.1.1. Ground Truth Flows. In order to obtain a baseline for assessing the performance of linearizations, of approximate numerical integration by various schemes or both, we always solved the assignment flow (nonlinear or linear) with high numerical accuracy using the geometric implicit Euler scheme (nonlinear flow) or the euclidean implicit Euler scheme (linear flow), with a sufficiently small step size h . This requires to solve a fixed point equation as part of every iteration k , involving the nonlinear mapping on the right-hand side of (3.10b) in

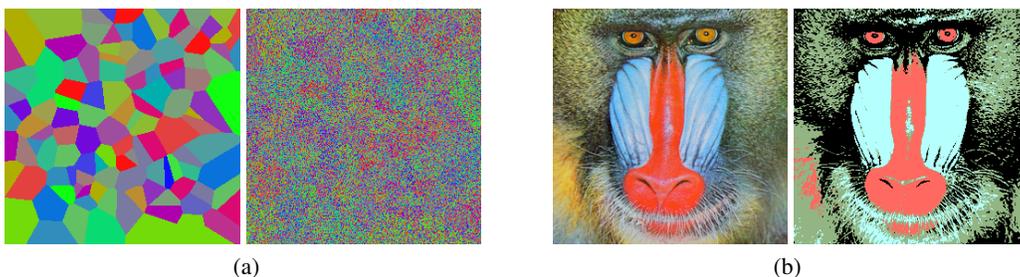


FIGURE 6.1. Labeling scenarios for evaluating numerical schemes for integrating the assignment flow. (a) Computer-generated data with 31 labels (left) and a noisy version used as input data (right). (b) A color image used as input data (left) using 4 color values as labels. These labels are illustrated on the right where each pixel has been replaced by the closest label.

case of the nonlinear assignment flow, or the linear mapping on the right-hand side of (4.3b) in case of the linear assignment flow. These fixed point equations were iteratively solved as well, and the corresponding iterations terminated when subsequent elements of the corresponding subsequences $(V^{k_i})_{i \geq 0}$ that measure the residual of the fixed point equation, satisfied

$$d_I(V^{k_{i+1}}, V^{k_i}) = \frac{1}{|J|} \max_{i \in I} \|V_i^{k_{i+1}} - V_i^{k_i}\| \leq 10^{-8}. \quad (6.1)$$

Starting these inner iterative loops with $V^{k_0} = V^{(k)}$ and terminating with $V^{k_{i,\text{end}}}$, we set $V^{(k+1)} = V^{k_{i,\text{end}}}$ and continued with the outer iteration $k + 1$.

6.1.2. *Termination criterion.* As suggested by [ÅPSS17], all iterative numerical schemes generating sequences $(W^{(k)})$ were terminated when the average entropy of the assignment vectors dropped below the threshold

$$-\frac{1}{|I||J|} \sum_{i \in I} \sum_{j \in J} W_{ij}^{(k)} \log W_{ij}^{(k)} < 10^{-3}. \quad (6.2)$$

If this happens, then – possibly up to a tiny subset of pixels $i \in I$ – all assignment vectors W_i are very close to a unit vector and hence almost uniquely indicate a labeling.

6.1.3. *Data.* Besides using the one-dimensional signal shown by Figure 6.4 that enables to visualize the entire evolution of the flow as plot in a single simplex (cf. Figure 6.6), we used two further labeling scenarios for evaluating the numerical integration schemes.

Figure 6.1(a) shows a scenario adopted from [ÅPSS17, Fig. 6] using $\rho = 0.1$ and $|\mathcal{N}_i| = 7 \times 7$. The input data (right panel) comprise 31 labels encoded as vertices (unit vectors) of a corresponding simplex. This results in *uniform* distances (2.17) and enables to assess in an unbiased way the effect of regularization by geometric diffusion in terms of the similarity map (2.19).

Figure 6.1(b) shows a color image together with 4 color vectors used as labels, as illustrated by the panel on the right. In contrast to the data of Figure 6.1(a) with a high level of noise and a uniform data term (as motivated and explained above), the input data shown on the left of Figure 6.1(b) are not noisy but comprise spatial structures at quite different scales (fine texture, large homogeneous regions), causing a nonuniform data term and a more complex assignment flow.

Both scenarios together provide a testbed in order to check and compare schemes for numerically integrating the assignment flow.

6.2. **Nonlinear Flow: Embedded RKMK-Schemes.** Figures 6.2 and 6.3 show the results of the two embedded RKMK schemes of Section 5.2 used to integrate the full nonlinear assignment flow (2.20), for the data shown by right panels of Figure 6.1 (a) and (b).

The two embedded RKMK schemes combine RKMK schemes of different approximation order q/q' , $1/2$ and $3/2$, respectively, which reuse vector field evaluations (3.17) in order to produce sequences of tangent vectors

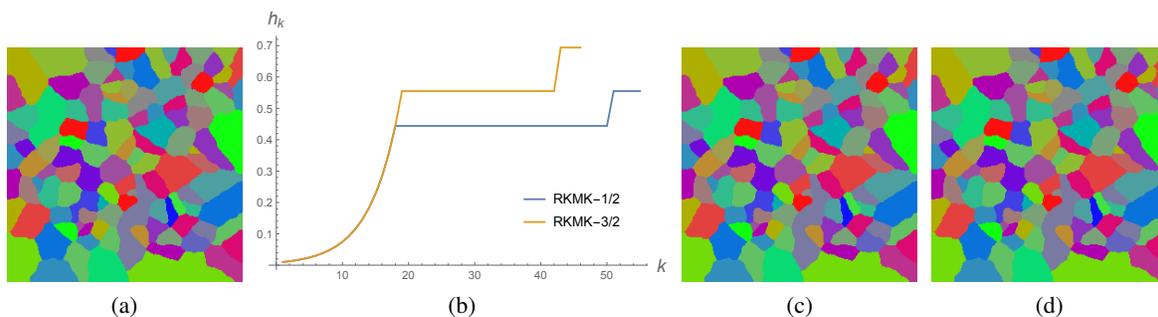


FIGURE 6.2. **Nonlinear assignment flow, embedded RKMK schemes.** Results of processing the data shown by Fig. 6.1(a), right panel (parameter: $\rho = 0.1$, $|\mathcal{N}_i| = 7 \times 7$). (a) Ground truth labeling resulting from integrating the (full nonlinear) assignment flow using the *implicit* geometric Euler scheme (step size $h = 0.5$). (b) Sequences (h_k) of adaptive step sizes generated by the embedded geometric schemes RKMK-1/2 and RKMK-3/2 of Section 5.2. The corresponding labeling results of these *explicit* schemes are shown as panels (c) and (d). Because there is almost no difference to the ground truth result, both explicit schemes integrate the assignment flow sufficiently accurate.

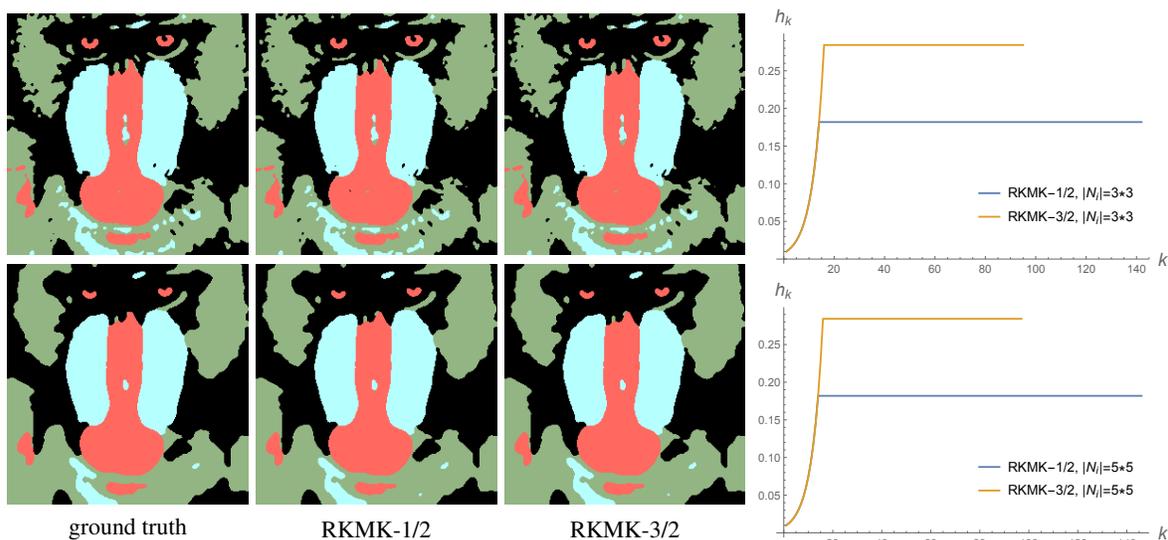


FIGURE 6.3. **Nonlinear assignment flow, embedded RKMK schemes.** Results of processing the data shown by Fig. 6.1(b), right panel (parameters: $\rho = 0.5$, $|\mathcal{N}_i| = 3 \times 3$ (less regularization; top row) and $|\mathcal{N}_i| = 5 \times 5$ (more regularization; bottom row)). The ground truth labeling was computed by integrating the (full nonlinear) assignment flow using the *implicit* geometric Euler scheme (step size $h = 0.5$). The embedded RKMK schemes (Section 5.2) generated the adaptive step sizes shown in the panels on the right and almost identical labeling results.

$(V^{(k)})$, $(\hat{V}^{(k)})$ that enable to estimate the local approximation error by monitoring the distances $d_I(V^{(k)}, \hat{V}^{(k)})$. As specified by (5.23), step sizes h_k adaptively increase provided a prescribed error tolerance is not violated.

The parameter values $\tau = 0.01$ (tolerance) and $n_\tau = 20$ (tolerance factor), used to produce the results shown by Figures 6.2 and 6.3, suffice to integrate accurately the full nonlinear assignment flow by the respective *explicit* schemes, as the comparison with the ground truth labeling generated by the implicit geometric Euler scheme shows. Since RKMK-3/2 has higher order q than RKMK-1/2, larger step sizes can be tolerated (panel (b)). On the other hand, each iteration of RKMK-3/2 is about twice expensive as RKMK-1/2.

Both plots of the step size sequences (h_k) reveal that the initial step size $h_0 = 0.01$ was much too small (conservative), and that a fixed value of h_k is adequate for most of the iterations. This value was larger for the experiment corresponding to Figure 6.2 due to the uniform data term (by construction, as explained in Section

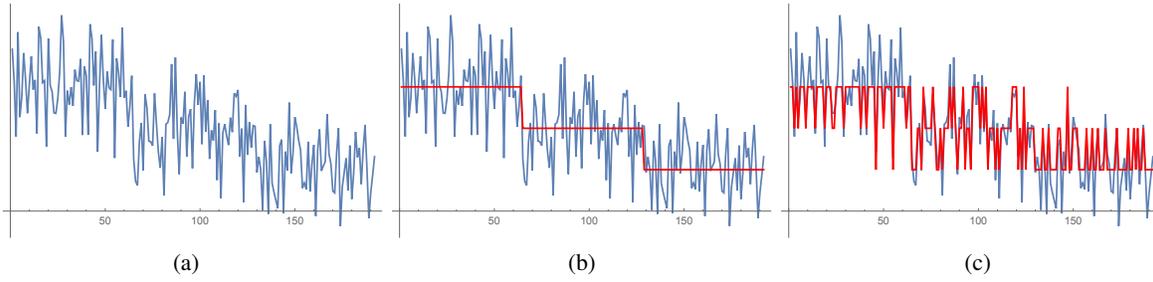


FIGURE 6.4. (a) A noisy 1D signal used for the experiments of Section 6.3.1. (b) The piecewise constant signal (red) used to generate (a) by superimposing noise. (c) *Local* rounding to the next label and comparing to (b) indicates the noise level. Local rounding is equivalent to omitting regularization in the assignment flow by replacing the interacting similarity vectors $S(W)$ in (2.20) by the non-interacting likelihood vectors $L(W)$ of (2.18).

6.1.3) and the more uniform scale of spatial structures. By contrast, the presence of spatial structures at quite different scales in the data corresponding to Figure 6.3 causes a more involved assignment flow to be integrated, and hence to a smaller step size after adaption. Comparing the two rightmost panels of Figure 6.3 shows that the strength of regularization (neighborhood size $|\mathcal{N}_i|$) had only little influence on the sequences of step sizes.

We never observed decreasing step sizes in these *supervised* scenarios, that is step 5. of (5.23) never was active. This may change in more involved scenarios, however (cf. Remark 3.2).

Overall, a few dozens of explicit iterations suffice for accurate geometric numerical integration of the assignment flow. Each iteration may be implemented in a fine-grained parallel way and has computational costs roughly equivalent to a convolution, besides mapping to the tangent space \mathcal{T}_0 and back to the assignment manifold \mathcal{W} , at each iteration.

6.3. Linear Assignment Flow. The approach of Section 4 involves two different approximations:

- (i) the *linear* assignment flow (4.2) approximating the *full* assignment flow (2.20), and
- (ii) the numerical integration of the linear assignment flow using two alternative numerical schemes:
 - (a) adaptive RK schemes (Section 5.1) based on the parametrization of Prop. 4.2 and
 - (b) the exponential integrator (Section 4.2).

Due to the remarkable approximation properties of the linear assignment flow when a *single* linearization at the barycenter is only used (Section 6.3.1), we entirely focused on this flow when evaluating the numerical schemes (a) and (b) in Sections 6.3.2 and 6.3.3.

6.3.1. Approximation Property. We report a series of experiments for the 1D signal depicted by Figure 6.4 using both the full and the linear assignment flow in order to check how closely the latter approximates the former. Then we discuss the linear assignment flow for the two 2D scenarios shown by Figure 6.1.

The parameter value $\rho = 0.1$ for scaling in (2.18) the data, for all 1D experiments discussed below. This gave a larger weight to the ‘data term’ so that – in view of the noisy data (Fig. 6.4) – the regularization property of the assignment flow (2.20), in terms of the similarity vectors $S_i(W)$ interacting through (2.19), was essential for labeling.

We first explain how the linearizations of the assignment flow were controlled. According to Proposition 4.2, using the parametrization (4.3a) and the linear ODE (4.3b) is equivalent to the linear assignment flow (4.2). Using again the parametrization (4.3a) and repeating the proof of Prop. 4.2 shows that the full assignment flow (2.20) is locally governed by the *nonlinear* ODE

$$\dot{V} = \Pi_{W_0}(S(\exp_{W_0}(V))), \quad V(0) = 0. \quad (6.3)$$

Taking into account (4.6) and subtracting the right-hand side of the approximation (4.3b) from the above right-hand side gives

$$\Pi_{W_0}(S(\exp_{W_0}(V))) - \Pi_{W_0}(s_0 + S_0 V) = \Pi_{W_0}(S(\exp_{W_0}(V)) - S(W_0) - dS_{W_0}(V)), \quad (6.4)$$

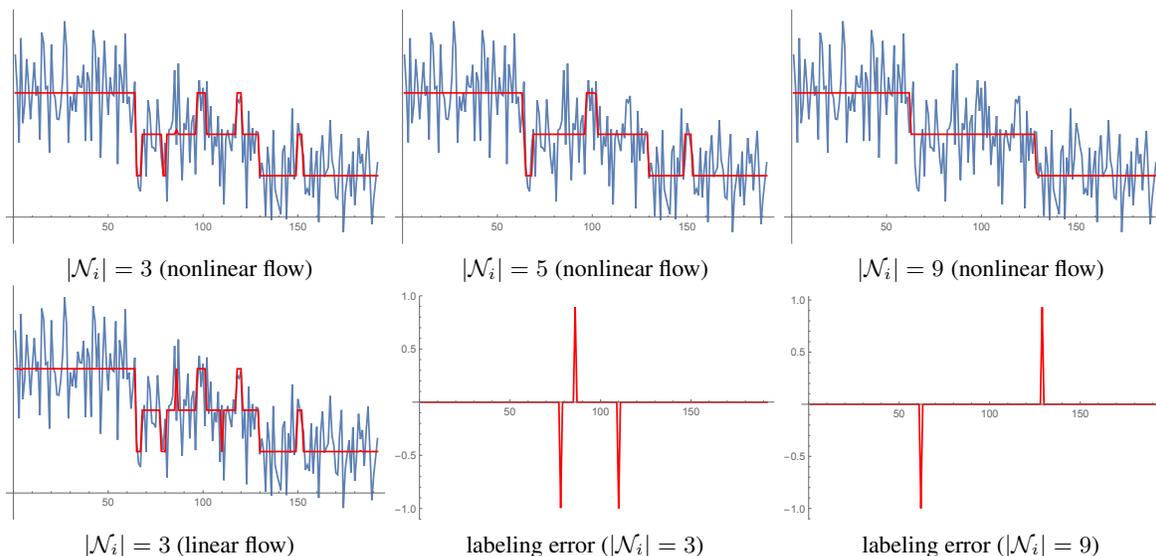


FIGURE 6.5. **Nonlinear vs. linear assignment flow.** TOP ROW: Labelings determined by the assignment flow as a reference for the linear assignment flow, using different neighborhood sizes $|\mathcal{N}_i|$. BOTTOM ROW: Labeling determined by the linear assignment flow that differs in 3 pixels from the corresponding above result of the full assignment flow. These errors and the two errors in the case $|\mathcal{N}_i| = 9$ are shown in the center and right panel. These and further results listed as Table 1 show that the linear assignment flow achieves high-quality labelings.

which shows that this approximation deteriorates with increasingly large tangent vectors V . As a consequence, we first solved the linear flow (4.2) using (4.3) *without* updating the point of linearization $W_0 = \mathbb{1}_{\mathcal{W}}$ and fixed after termination at k_{end} (= number of required outer iterations) the constant

$$\|V\|_{\max} = \max_{i \in I} \|V_i^{(k_{\text{end}})}\|. \quad (6.5)$$

Then we solved the linear assignment flow again and updated the linearization point W_0 in view of (6.4) whenever

$$\max_{i \in I} \|V_i^{(k)}\| > \frac{\|V\|_{\max}}{c}, \quad c \geq 1, \quad (6.6)$$

using the parameter c to control the number of linearizations: a single linearization and no linearization update if $c = 1$ and an increasing number of updates for larger values of c . We updated components of the linearization point $W_{0,i}$ by $W_i(h)$, $i \in I$ only when $\min_{j \in J} W_{i,j}(h) > 0.01$, in order to keep linearization points inside the simplex, in view of the entries (4.7d) of dS_{W_0} normalized by components of W_{0k} .

After termination, the induced labelings were compared to those of the full assignment flow, and the number of wrongly assigned labels was taken as a quantitative measure for the approximation property of the linear assignment flow: Except for the minimal neighborhood size $|\mathcal{N}_i| = 3$, a single linearization almost suffices to

$ \mathcal{N}_i $	$c = 1$	$c = 2$	$c = 3$	$c = 4$	$c = 5$
3	1/3	4/3	7/3	10/1	13/0
5	1/0				
9	1/2	5/0			

TABLE 1. Approximation property of the linear assignment flow. The entries x/y specify the number x of linearizations and the number y of wrongly assigned labels (out of 192 assigned labels), depending on the neighborhood size $|\mathcal{N}_i|$ (strength of regularization) and the parameter c specifying the tangent space threshold (6.6).

obtain a correct labeling. Overall, the maximal number of 3 labeling errors (out of 192) is very small, and these

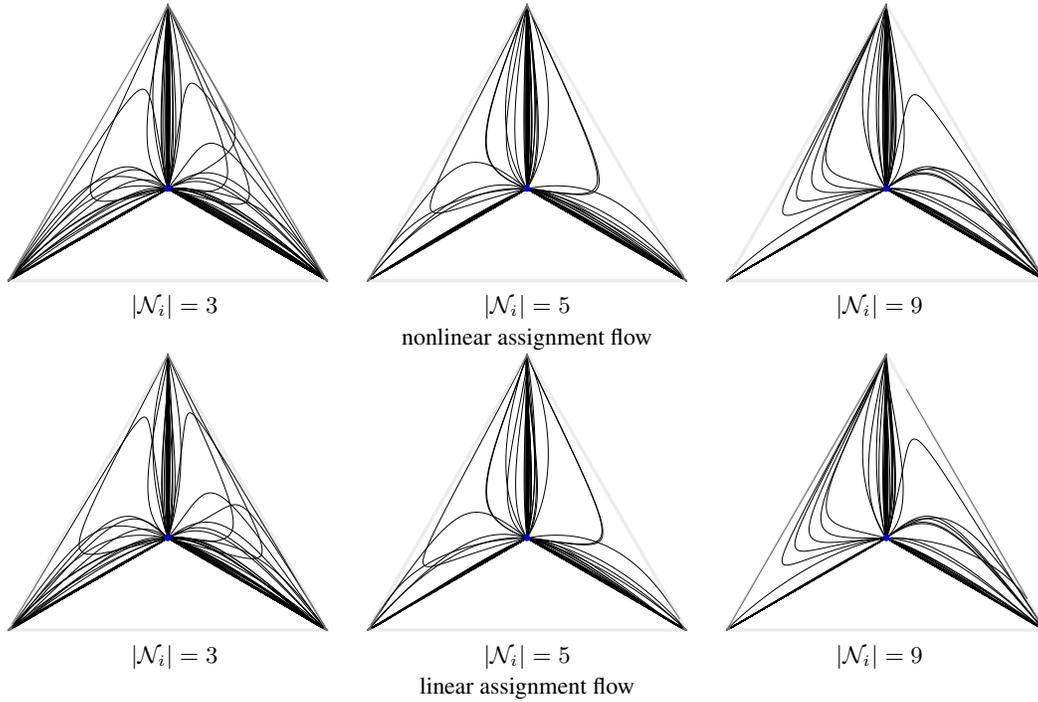


FIGURE 6.6. **Nonlinear vs. linear assignment flow.** Comparison of the assignment flow (2.20) (top row) and the linear assignment flow (4.2) (bottom row) in terms of all $|I|$ solution curves $W_i(t)$, $i \in I$, plotted in the 3-label simplex. A major part of the trajectories approaches more or less directly a vertex, whereas another part changes the original direction due to regularization by geometric smoothing. Except for the cases with the minimal neighborhood $|\mathcal{N}_i| = 3$, the similarity of both flows is apparent. This illustrates the reason for very small observed numbers of labeling errors, as listed and depicted by Table 1 and Figure 6.5.



FIGURE 6.7. **Linear assignment flow.** Labeling results for the two scenarios of Figure 6.7, using the *linear* assignment flow with a single linearization at the barycenter and the *implicit* Euler scheme for numerical integration. Comparison with the labeling results of the *nonlinear* assignment flow (Fig. 6.2(a), Fig. 6.3 ‘ground truth’) demonstrates a remarkable approximation property of the linear assignment flow.

errors merely correspond to shifts by a single pixel position of the signal transition in the case $|\mathcal{N}_i| = 9$ (see Figure 6.5). We conclude that for supervised labeling, the linear assignment flow (4.2) (which is nonlinear!) – cf. Remark 4.3) indeed captures a major part of nonlinearity of the full assignment flow (2.20). Figure 6.6 illustrates the similarity of the two flows (and the dissimilarity in the case $|\mathcal{N}_i| = 3$) in terms of all $|I| = 192$ sequences $(W_i^{(k)})$, $i \in |I|$, plotted as piecewise linear trajectories.

We cannot assure, however, that this approximation property persists in more general cases (cf. Remark 3.2) whose study is beyond the scope of the present paper.

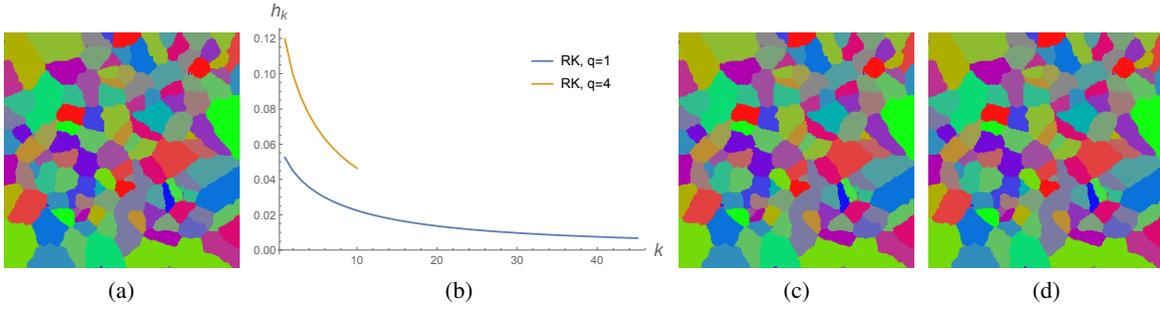


FIGURE 6.8. **Linear assignment flow, adaptive RK-schemes.** Results of the linear assignment flow (4.2) based on the parametrization (4.3), the RK schemes (5.1) of order $q = 1$ (FE) and $q = 4$ (RK4), and adaptive step size selection based on the local error estimate (5.8). The labeling results (c), (d) for $q = 1$, $q = 4$ are almost identical to ground truth (a) from Figure 6.7.

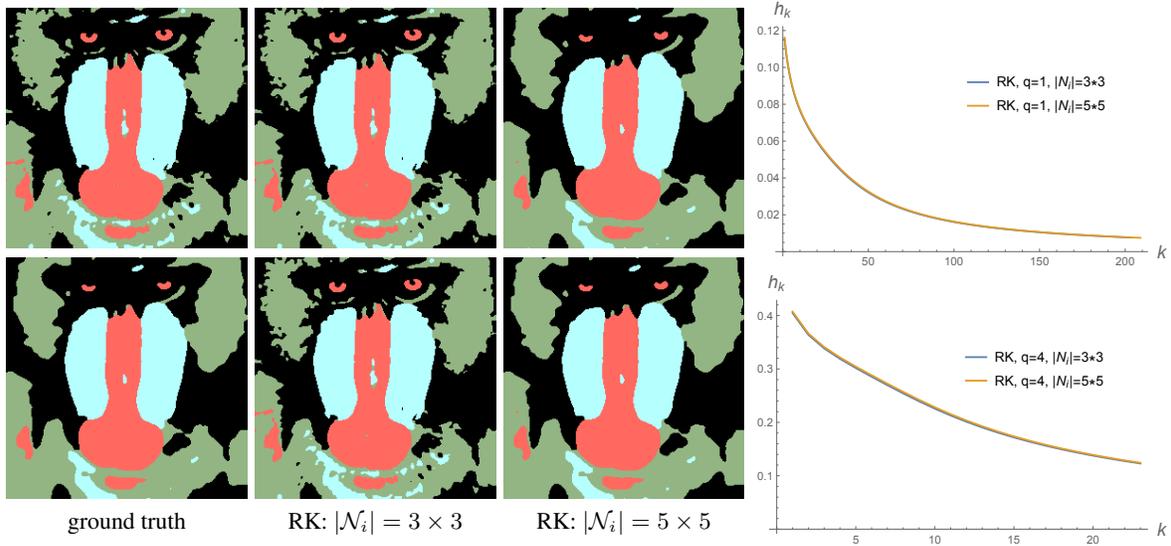


FIGURE 6.9. **Linear assignment flow, adaptive RK-schemes.** Results of the linear assignment flow (4.2) based on the parametrization (4.3), the RK schemes (5.1) of order $q = 1$ (FE) and $q = 4$ (RK4), and adaptive step size selection based on the local error estimate (5.8). For fixed q , increasing the neighborhood size $|\mathcal{N}_i|$ has almost no effect: the step size sequences agree up to the second digit. The labeling results (c), (d) for $q = 1$, $q = 4$ are almost identical to ground truth from Figure 6.7.

We now turn to the scenarios shown by 6.1. Figure 6.7 shows the results obtained using the implicit Euler scheme and the *same* parameter settings that were used to integrate the nonlinear flow, to obtain the ground truth flows and results depicted by Figure 6.2(a) and Figure 6.3, left-most panel, respectively. Comparing the labelings returned by the linear and nonlinear assignment flow, respectively, confirms the discussion of the 1D experiments detailed above: The results agree except for a very small subset of pixels close to signal transitions which are immaterial for subsequent image interpretation.

The results shown by Figure 6.7 served as ground truth for studying the *explicit* numerical schemes of Sections 6.3.2 and 6.3.3 for integrating the linear assignment flow.

6.3.2. Adaptive RK Schemes. We evaluated the adaptive RK schemes (FE) of order $q = 1$ and (RK4) of order $q = 4$, due to (5.1), supposed to integrate the linear ODE (4.3b), after rearranging the polynomials of (5.1) in Horner form.

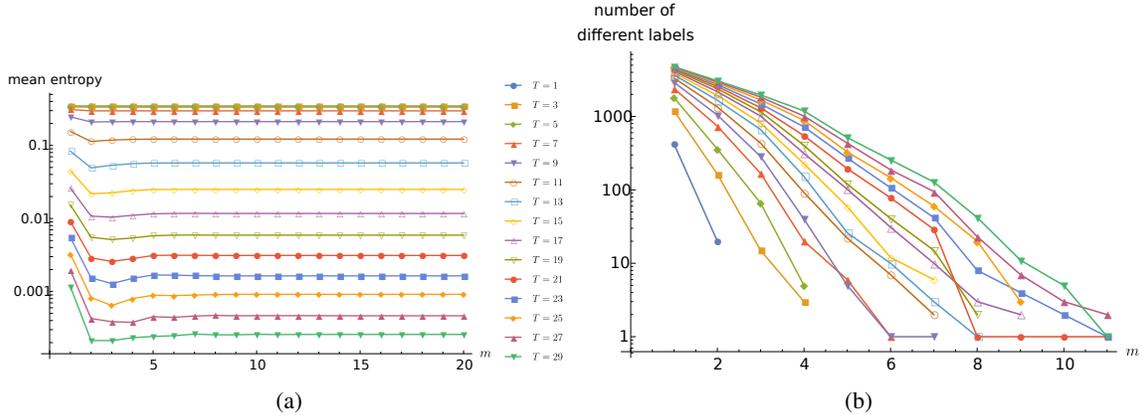


FIGURE 6.10. **Linear assignment flow, exponential integrators, influence of the Krylov subspace dimension m and the point of evaluation T .** The results correspond to the labeling problem shown by Figure 6.11. They demonstrate that m can be chosen quite small. Figure (a) shows that the mean entropy decreases with increasing T but does not decrease for fixed T and $m > 3$. Figure (b) shows for $T = 1, \dots, 29$ (same color code as (a)) and for each m the number of labels that change when the dimension of the Krylov subspace is increased to $m + 1$. All curves decrease with increasing m , and curves that reach 0 label changes just discontinue due to the logarithmic scale. The plot shows that independent of T (i.e. for any T) there is an m such that increasing the dimension of the Krylov space does not change the labeling at all.

Figures 6.8 and 6.9 show the results for the *linear* assignment flow based on a single linearization at the barycenter, using the results shown by Figure 6.7 as ground truth. The step sizes h_k were computed at each iteration k using the local error estimate (5.8) such that $\frac{1}{|I|^{1/2}} \|V(t_{k+1}) - V^{(k+1)}\| \leq \tau = 0.01$, that is on average $\|V_i(t_{k+1}) - V_i^{(k+1)}\| \leq \tau$ for all pixels $i \in I$. The spectral norm $\|A\|$ was computed beforehand using the basic power iteration.

As explained above when the criterion (5.19) was introduced, step sizes must decrease due to the increasing norms $\|V^{(k)}\|$, in order to keep bounded the local integration error. Furthermore, in agreement with Figure 5.1, raising the order q of the integration scheme leads to significantly larger step sizes and hence to smaller total numbers of iterations, at the cost of more expensive iterative steps. Yet, roughly taking these additional costs into account by multiplying the total iteration numbers for $q = 4$ by 4, indicates that raising the order q reduces the overall costs. In this respect our findings for the linear assignment flow differ from the corresponding findings for the nonlinear assignment flow and the embedded RKMK schemes, discussed in Section 6.2.

6.3.3. Exponential Integrator. For integrating the linearized assignment flow with exponential integrators, we consider equation (4.19) and the Krylov space approximation (4.23)

$$V(T) = T\varphi_1(TA)a \approx T\|a\|V_m\varphi_1(TH_m)e_1. \quad (6.7)$$

As the evaluation of $\varphi_1(TH_m)e_1$ is explained in Section 4.2, we only discuss here the choice of the parameters m and T .

The dimension m of the Krylov subspace controls the quality of the approximation, where larger values theoretically lead to a better approximation. In our experiments, rather small numbers, like $m = 5$, turned out to suffice to produce labelings very close to the ground truth labelings, that were generated by the implicit Euler method – see Figures 6.10 and 6.11. As the runtime of the algorithm increases with growing m , this parameter should not be chosen too large.

Scaling A and a in (4.18) affects the vector field defining the linear ODE (4.3b). Hence, fixing any time point T depends on this scaling factor, too. As a consequence, since A and a depend on the problem data (4.18), the choice of T is problem dependent. On the other hand, the discussion following the proof of Theorem 5.2 showed that $\|V(t)\|$ increases with t , and T merely has to be chosen large enough such that $W(T)$ defined by

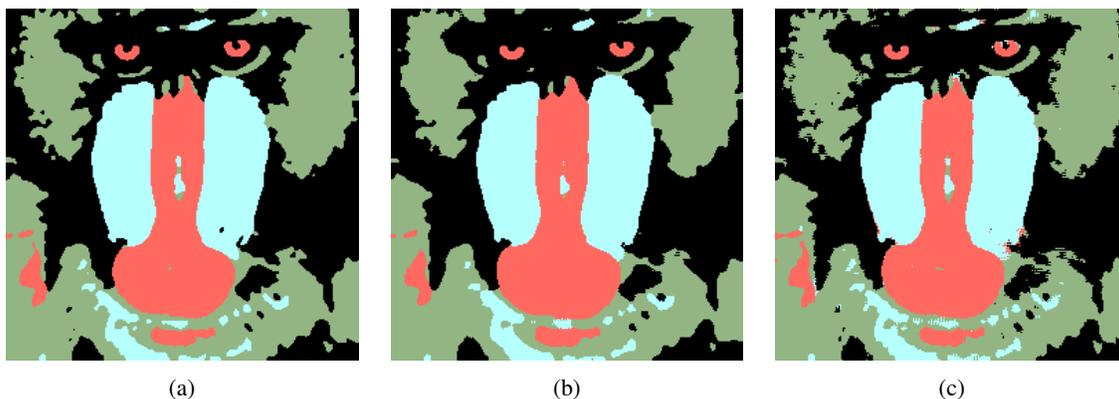


FIGURE 6.11. **Linear assignment flow, exponential integrators, sample labelings.** Figure (a) shows the ground truth labeling as generated by the implicit Euler method. Figure (b) displays the labeling generated by the exponential integrator using the Krylov subspace dimension $m = 5$, which is very close to the ground truth labeling. As demonstrated by Figure 6.10, m cannot be chosen too small, however, since labelings then start to deteriorate rapidly. This is illustrated by Figure (c) which shows the labeling for $m = 3$. Comparison with (b) shows that dimensions 4 and 5 ‘contain’ small-scale details of the correct labeling induced by the linear assignment flow.

(4.3a) satisfies the termination criterion (6.2) – see (5.21c) for a rough estimate. Choosing T overly large will cause numerical underflow and overflow issues, however.

Almost all runtime is consumed by the Arnoldi iteration producing the subspace basis V_1, \dots, V_m . Due to the small dimension m , the total runtime is very short, and time required for the subsequent evaluation of the right-hand side of (6.7) is neglectable.

7. CONCLUSION

We investigated numerical methods for image labeling by integrating the large system of nonlinear ODEs defining the assignment flow (2.20), which evolves on the assignment manifold. All methods exactly respect the underlying geometry. Specifically, we adapted RKMK methods and showed that *embedded* RKMK-methods work very well for automatically adjusting the step size, at neglectable additional costs. Raising the order enables leads to larger step sizes, which is compensated by the higher computational costs per iteration, however. In either case, each iteration only involves convolution like operations over local neighborhoods together with pixelwise nonlinear functions evaluations.

We derived and introduced the *linear* assignment flow, a nonlinear approximation of the (full) assignment flow that is governed by a large linear system of ODEs on the tangent space. Experiments showed that the approximation is remarkably close, as measured by the number of different label assignments.

We investigated two further families of numerical schemes for integrating the linear assignment flow: established RK schemes with adaptive step size selection based on a local integration error estimate, and exponential integrators for approximately evaluating Duhamel’s integral using a Krylov subspace. In the former case, higher-order schemes really pay, unlike for the RKMK schemes and the full assignment flow, as mentioned above. Choosing the classical RK scheme with $q = 4$, for example, few dozens of iterations suffice to reach the termination criterion, with high potential for parallel implementation. The exponential integrators, on the other hand, directly approximate the integral determining $V(T)$ and in this sense are non-iterative. Here, a Krylov subspace basis of low dimension merely has to be computed, using a standard iterative method. Even though this method differs mathematically from the RK schemes, it has potential for real-time implementation as well.

All methods provide a sound basis for more advanced image analysis tasks that involve labeling by evaluating the assignment flow as a subroutine. Accordingly, our future work concerns an extension of the unsupervised labeling approach [ZZr⁺18], where label dictionaries are directly learned from data through label assignment. Furthermore, methods under investigation for learning to adapt regularization parameters of the assignment flow

to specific image classes, require advanced discretization and numerical methods based on the results reported in the present paper.

Acknowledgements. This work was supported by the German Research Foundation (DFG), grant GRK 1653.

REFERENCES

- [AE05] N. Ay and I. Erb, *On a Notion of Linear Replicator Equations*, J. Dynamics Diff. Eq. **17** (2005), no. 2, 427–451.
- [AJLS17] N. Ay, J. Jost, H. V. Lê, and L. Schwachhöfer, *Information Geometry*, Springer, 2017.
- [Alz97] H. Alzer, *On Some Inequalities for the Incomplete Gamma Function*, Math. Comp. **218** (1997), no. 66, 771–778.
- [AMH11] A. Al-Mohy and N. Higham, *Computing the Action of the Matrix Exponential, with an Application to Exponential Integrators*, SIAM Journal on Scientific Computing **33** (2011), no. 2, 488–511.
- [AN00] S.-I. Amari and H. Nagaoka, *Methods of Information Geometry*, Amer. Math. Soc. and Oxford Univ. Press, 2000.
- [ÅPSS17] F. Åström, S. Petra, B. Schmitzer, and C. Schnörr, *Image Labeling by Assignment*, Journal of Mathematical Imaging and Vision **58** (2017), no. 2, 211–238.
- [BFPS17] R. Bergmann, J. H. Fitschen, J. Persch, and G. Steidl, *Iterative Multiplicative Filters for Data Labeling*, Int. J. Computer Vision **123** (2017), no. 3, 435–453.
- [E17] Weinan E, *A Proposal on Machine Learning via Dynamical Systems*, Communications in Mathematics and Statistics **5** (2017), no. 1, 1–11.
- [Hig08] N.J. Higham, *Functions of Matrices: Theory and Computation*, SIAM, 2008.
- [HL97] M. Hochbruck and C. Lubich, *On Krylov Subspace Approximations to the Matrix Exponential Operator*, SIAM Journal on Numerical Analysis **34** (1997), no. 5, 1911–1925.
- [HLW06] E. Hairer, C. Lubich, and G. Wanner, *Geometric Numerical Integration*, Springer, 2006.
- [HNW08] E. Hairer, S.P. Nørsett, and G. Wanner, *Solving Ordinary Differential Equations I*, 3rd ed., Springer, 2008.
- [HO10] M. Hochbruck and A. Ostermann, *Exponential Integrators*, Acta Numerica **19** (2010), 209–286.
- [HR17] E. Haber and L. Ruthotto, *Stable Architectures for Deep Neural Networks*, Inverse Problems **34** (2017), no. 1, 014004.
- [HSÅS18] R. Hühnerbein, F. Savarino, F. Åström, and C. Schnörr, *Image Labeling Based on Graphical Models Using Wasserstein Messages and Geometric Assignment*, SIAM J. Imaging Science **11** (2018), no. 2, 1317–1362.
- [HZRS16] K. He, X. Zhang, S. Ren, and J. Sun, *Deep Residual Learning for Image Recognition*, Proc. CVPR, 2016.
- [IMKNZ05] A. Iserles, H. Z. Munthe-Kaas, S. P. Nørsett, and A. Zanna, *Lie-Group Methods*, Acta Numerica **14** (2005), 1–148.
- [KAH⁺15] J.H. Kappes, B. Andres, F.A. Hamprecht, C. Schnörr, S. Nowozin, D. Batra, S. Kim, B.X. Kausler, T. Kröger, J. Lellmann, N. Komodakis, B. Savchynskyy, and C. Rother, *A Comparative Study of Modern Inference Techniques for Structured Discrete Energy Minimization Problems*, International Journal of Computer Vision **115** (2015), no. 2, 155–184.
- [MK99] H. Munthe-Kaas, *High order Runge-Kutta Methods on Manifolds*, Applied Numerical Mathematics **29** (1999), no. 1, 115–127.
- [MVL03] C. Moler and C. Van Loan, *Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later*, SIAM Review **45** (2003), no. 1, 3–49.
- [NW12] J. Niesen and W.M. Wright, *Algorithm 919: A Krylov Subspace Algorithm for Evaluating the φ -Functions Appearing in Exponential Integrators*, ACM Trans. Math. Software **38** (2012), no. 3, Article 22.
- [Saa92] Y. Saad, *Analysis of Some Krylov Subspace Approximations to the Matrix Exponential Operator*, SIAM Journal on Numerical Analysis **29** (1992), no. 1, 209–228.
- [Sid98] R.B. Sidje, *Expokit: A Software Package for Computing Matrix Exponentials*, ACM Transactions on Mathematical Software (TOMS) **24** (1998), no. 1, 130–156.
- [Tes12] G. Teschl, *Ordinary Differential Equations and Dynamical Systems*, Grad. Studies Math., vol. 140, Amer. Math. Soc., 2012.
- [ZZr⁺18] A. Zern, M. Zisler, F. Åström, S. Petra, and C. Schnörr, *Unsupervised Label Learning on Manifolds by Spatially Regularized Geometric Assignment*, Proc. GCPR, 2018.

(A. Zeilmann) IMAGE AND PATTERN ANALYSIS GROUP, HEIDELBERG UNIVERSITY, GERMANY
E-mail address: alexander.zeilmann@iwr.uni-heidelberg.de

(F. Savarino) IMAGE AND PATTERN ANALYSIS GROUP, HEIDELBERG UNIVERSITY, GERMANY
E-mail address: fabrizio.savarino@iwr.uni-heidelberg.de

(S. Petra) MATHEMATICAL IMAGING GROUP, HEIDELBERG UNIVERSITY, GERMANY
E-mail address: petra@math.uni-heidelberg.de
URL: <https://www.stpetra.com>

(C. Schnörr) IMAGE AND PATTERN ANALYSIS GROUP, HEIDELBERG UNIVERSITY, GERMANY
E-mail address: schnoerr@math.uni-heidelberg.de
URL: <https://ipa.math.uni-heidelberg.de>