

A Differential-Algebraic Projective Framework for the Deformable Single-View Geometry of the 1D Perspective Camera

Adrien Bartoli

ENCOV, Institut Pascal

UMR6602 CNRS, Université Clermont Auvergne, SIGMA

Corresponding author: Adrien Bartoli

`Adrien.Bartoli@gmail.com`

Abstract

Single-View Geometry (SVG) studies the world-to-image-mapping or warp, which is the relationship that exists between a body's model and its image. For a rigid body observed by a projective camera, the warp is described by the usual camera matrix and its properties. However, it is clear that for a body whose deformation state changes between the body's model and its image, the 'simple', globally parameterized warp described solely by the camera matrix breaks down. Existing work has exploited deformation to reconstruct the deformed body from its image, but did not establish the properties of the deformable warp. Studying these properties is part of deformable SVG and forms a recent research topic. Because deformations may take place anywhere in the object's body, and because they may be uncorrelated, the warp is local in nature. Using a differential framework is thus an obvious choice. We propose a differential-algebraic projective framework based on modeling the body's surface by a locally-rational projective embedding and on the 1D projective camera. We show that this leads, via the study of univariate rational functions, to differential invariants that the warp must satisfy. It may seem surprising, given the generic hypothesis made on the observed body, hardly stronger than mere local-smoothness, that constraints can still be found. Our framework generalizes the Schwarzian derivative, the first-order projective differential invariant, which holds under the assumption that the body's shape is locally-linear. Our invariants may be used to construct regularizers to be used in warp estimation. We report experimental results of two types on simulated and real data. The first type shows that the proposed invariants hold well for an independently estimated warp. The second type shows that the proposed regularizers improve warp estimation from point correspondences compared to the classical derivative-penalizing regularizers.

1 Introduction

Single-View Geometry (SVG) is concerned with how cameras form images of the world geometrically. More precisely, SVG deals with the geometry and algebra of the *warp*, a mapping that exists between a model of the world and a model of the image, as illustrated by figure 1. SVG endeavors to (i) establish the existence and form of such mappings, (ii) study their characteristics and (iii) estimate them numerically. For rigid bodies, SVG uses simple geometric primitives such as points and an algebraic projective framework, which led to a deep understanding of points (i), (ii) and (iii) (Faugeras, 1993; Forsyth and Ponce, 2012; Hartley and Zisserman, 2003). Point (i) merely uses the body’s pose and the camera’s pin-hole projection model, and leads to the projective camera represented by a (3×4) matrix, point (ii) studies what properties make a (3×4) matrix valid as a camera and point (iii) solves the so-called camera resection or pose estimation problem. The case of deformable bodies forms an open and challenging research challenge, owing to the deformation causing the model and observed body shapes to be different. Deformable SVG is important because it will form the mathematical framework for and deepen our understanding of the Shape-from-Template (SfT) problem (Gumerov et al., 2006; Perriollat et al., 2011; Salzmann et al., 2007). In order to define deformable SVG more specifically, we start with the informal definition of an important mapping, the *embedding*. The embedding maps points from the body’s model to its deformed state in camera coordinates, and thus extends the notion of relative body-camera pose from the rigid case as it also holds the body’s deformation. As in the rigid case, the *projection* mapping represents the camera, mapping points from the deformed body to its observed image. The warp is then defined as the composition of the projection with the embedding. It maps points from the body’s model to the observed image of the deformed body. This informal definition of the warp already provides a simple and generic answer to point (i). In contrast, point (ii) is much more involved, but tremendously important to understand, for two main reasons. The first reason is to understand the theoretical possibility of characterizing the observations forming valid images *independently* of the body’s deformed state. It contrasts with the rigid case where the body’s state is already known: in the deformable case, the observed shape is unknown because deformation occurred between the body’s state in the model and the body’s state as observed in the image. The second reason is that the theoretical characteristics of the warp may be used to improve point (iii). For deformable bodies, it is indeed very common to estimate the warp using a derivative-penalizing regularizer, as in the Thin-Plate Spline (Bookstein, 1989) and optic flow computation (Fortun et al., 2015). An understanding of point (ii) can provide physically sound regularizers, with the potential to improve warp estimation. Importantly, these regularizers are independent of the observed shape and thus insensitive to reconstruction ambiguities. These considerations however require further assumptions on the embedding, for, by definition, the warp is so far

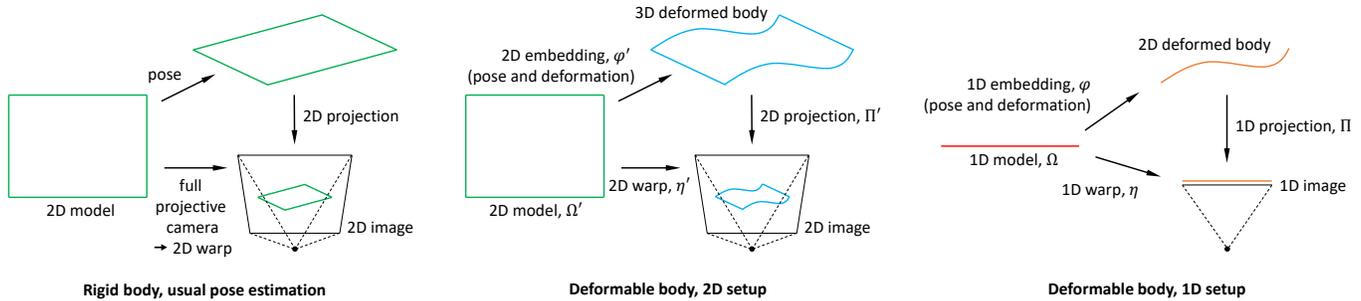


Figure 1: Geometric setup used in studying SVG. The mapping to be studied is the warp, which exists between the body’s model and the body’s image. In the rigid case, the image shows the body’s model up to pose and perspective projection. In the deformable case, the body’s shape changes between the model and image with much higher complexity and the setup uses the *embedding-projection* framework. The 1D setup is simpler to study than the 2D setup, yet provides valuable theoretical and practical insights for both cases.

just a mere locally-smooth mapping.

We propose to study deformable SVG for *locally-rational projective embeddings*. Locally-rational means that the embeddings resemble a rational function of some degree in the infinitesimal vicinity of any smooth point. A rational function is formed by the ratio of two polynomials and thus naturally includes simple polynomial functions. It is simple to see that the locally-rational embeddings lead to locally-rational warps, as the projection mapping is essentially a ratio. A noteworthy consequence is that the polynomial embeddings do not lead to a simplified form of the warp and thus of the SVG. Our motivations for choosing these embeddings are three-fold. Our first motivation is genericity: the chosen embeddings can represent the surface of thin or thick objects, whether they deform isometrically or in more complex ways. Indeed, any embedding can be locally approximated to a good extent by its Taylor or Padé approximation (Press et al., 2007) at some finite order. Understanding the limits of using such an extremely generic prior, slightly stronger than mere local-smoothness, forms a fundamental research question. A locally-linear embedding was successfully used in isometric deformable reconstruction (Parashar et al., 2018). The chosen embeddings generalize this concept, as they can behave locally as polynomials of a fixed but arbitrary degree, beyond linearity. Our experimental results show that in terms of warp estimation, the quadratic and cubic locally-rational models outperform the linear model, and so that the proposed generalization brings a practical improvement. Our second motivation is simplicity and stability of the representation. A global embedding model would have to adjust its complexity to the actual deformation. Therefore, coping with complex deformations would come at the price of a higher degree, potentially causing instabilities. In contrast, the chosen embeddings are locally constrained to behave similarly to rational functions of low degree. Our third motivation is that this representation allows us to truly model perspective projection without ever approximating it. Indeed, perspective projection is simply expressed by rationality, which is perfectly respected by the proposed

invariants.

Our methodology is to first study point *(ii)* for strictly rational embeddings, establishing differential invariants on the warp. Because the existence and study of such invariants for an arbitrary degree were not previously established, except for the special case of linear rational functions (Ovsienko and Tabachnikov, 2005), this forms a main contribution of our work. We specifically studied the differential invariants though other non-differential invariants may exist. We then study point *(iii)*, deriving warp regularizers from our differential invariants. We use these regularizers as soft penalties in warp estimation in place of the usual derivative-penalizing regularizers. Concretely, we have chosen to work with a 1D setup, for which the deformable body is a plane curve, the camera a projection from 2D to 1D and the warp a 1D mapping. Our motivations are two-fold. Our first motivation is that the 1D warp case forms a necessary step to establish the 2D warp case, where the deformable body is a surface and the camera a usual pin-hole. The 1D case is simpler than the 2D case, yet conveys valuable insights on the theoretical form of the sought invariants, their use in warp estimation and their practical impact. Indeed, our experimental results show that in terms of warp estimation, the locally-rational representation substantially outperforms the locally-polynomial representation assumed by the existing derivative-penalizing regularizers, when used in a compound cost exploiting point correspondences. Our second motivation is that the 1D warp invariants already give a subset of the 2D warp invariants. We show this fact by constructing a virtual 1D setup from the 2D setup. The main result is that the 1D invariants hold for a 1D warp along any rational curve chosen in the 2D model and for any 1D projection of the 2D image.

We first review previous work in §2 and give background material in §3. We then give our theoretical results on the invariants of univariate rational functions in §4 and show how this applies to deformable 1D SVG in §5. We give experimental results on warp estimation in §6 and a conclusion in §7. Finally, a first appendix discusses the extension of our 1D framework to higher dimensions, in particular to form a basis for 2D deformable SVG, and a second appendix gives the proof of our proposition regarding the invariants.

2 Previous Work

We split our review of previous work into four parts: deformable SVG, the 1D camera, rational functions and the embedding-projection framework.

Deformable SVG. SVG has been well studied for rigid bodies (Faugeras, 1993; Forsyth and Ponce, 2012; Hartley and Zisserman, 2003) and very scarcely for deformable bodies. Most of the work in SFT, whose geometric setup is modeled by SVG, focuses on computing the body’s 3D deformation and not on characterizing the warp’s properties (Gumerov et al., 2006; Perriollat et al., 2011; Salzmann et al., 2007). The estima-

tion of rigid SVG and Multiple-View Geometry (MVG) from images of algebraic curves was thoroughly studied (Kaminski and Shashua, 2004; Schmid and Zisserman, 2000). The deformable case of multiview reconstruction was specifically studied for isometric curves (Faugeras, 1993). A comprehensive differential framework was given in (Fabbri and Kimia, 2016). There are two main differences between these works and ours. First, they use a rigid or an isometric prior, while we use a locally-rational model, which is a much more generic, and thus much weaker constraint. Second, they study the differential relationships between space and image curves for quantities such as speed, curvature and torsion, while we establish generic invariants at any order.

The 1D camera. We use the 1D camera as a simplified case of the regular 2D camera. The 1D camera was introduced in (Quan and Kanade, 1997) for line-based affine Structure-from-Motion but can be also derived by analogy to the usual 2D camera. It has a total of 5 parameters which are 2 intrinsics (the focal length f and the principal point q_0 , both expressed in number of pixels) and 3 extrinsics (a rotation angle and two translation parameters). The strategy of using the 1D camera to simplify a 2D problem was used in the rigid case, for instance the 1D trifocal tensor is much easier to understand and work with than the 2D trifocal tensor (Faugeras and Mourrain, 1995; Faugeras et al., 2000). We defined the 1DSfT problem, and showed that it makes expressing a constraint such as isometry much easier than in the usual 2D formulation (Gallardo et al., 2015). We showed that, in general configuration, 1DSfT has a discrete number of solution if the isometric prior is used, but an infinite number of ambiguities otherwise. Solutions derived in the 1D case can sometimes be used directly in the 2D case. For instance, the geometry of a 2D camera in planar motion is equivalent to a 1D camera on the trifocal plane of the 2D cameras (Faugeras et al., 2000).

Rational functions. Rational functions are used in several areas of applied mathematics such as data interpolation via the Padé approximation (Press et al., 2007) and the Non-Uniform Rational B-Spline (NURBS) (Piegl and Tiller, 1997). Their properties were studied for instance in (Sendra et al., 2008; Walker, 1978). The closest fundamental result to ours is probably the Schwarzian derivative (Ovsienko and Tabachnikov, 2009), a differential version of the cross-ratio which we showed may be directly used in warp estimation (Pizarro et al., 2016). The Schwarzian is also the solution of the ODE $2\mu^{(1)}\mu^{(3)} - 3\mu^{(2)^2} = 0$ (Ovsienko and Tabachnikov, 2005), where $\mu^{(k)}$ represents the k -th derivative of μ . In the context of SVG it is derived by considering a linear projective embedding, or equivalently, making the assumption that the curve is locally flat. Our theory generalizes the Schwarzian by considering rational projective embeddings of an arbitrary degree and has the Schwarzian as a special case.

The embedding-projection framework. The embedding-projection framework has been extensively

used to model SfT. In the 2D setup, differential solutions were found to resolve reconstruction for the isometric and conformal models (Bartoli et al., 2015). The framework has more recently been used in the 1D setup to study isometric reconstruction (Gallardo et al., 2015). These works and their extensions use differential equations giving the reconstruction in terms of the warp. However, none of them studied deformable SVG. They thus assume that the warp is computed in a preliminary step, independently of the reconstruction, thereby neglecting valuable constraints. In contrast, our study of deformable SVG shows that the warp strongly depends on the assumptions made on the embedding. Deriving these warp invariants by marginalising the embedding is far from trivial, but using them in warp estimation then brings substantial improvements.

3 Background

3.1 Notation

We write logical equivalence as \Leftrightarrow . We use $C^\infty(\mathbb{R}^c, \mathbb{R}^d)$ for the set of smooth functions from \mathbb{R}^c to \mathbb{R}^d and use the shortcut $C^\infty \stackrel{\text{def}}{=} C^\infty(\mathbb{R}, \mathbb{R})$. We write $\mathcal{U} = [a, b]$ the set of natural numbers in the interval $[a, b]$, for $a, b \in \mathbb{N}$. For a set $\mathcal{U} \subset \mathbb{N}$, $\mathcal{U} = [u_1, u_{|\mathcal{U}|}]$, we define $\mu^{(\mathcal{U})} \stackrel{\text{def}}{=} (\mu^{(u_1)} \dots \mu^{(u_{|\mathcal{U}|})})^\top$ as a multivalued function giving the orders of derivatives of μ from \mathcal{U} . We use \mathbb{R}^* to denote the set of non-zero real numbers and \mathbb{N}^* to denote the set of non-zero natural numbers.

3.2 Univariate Polynomial Functions

Let \mathcal{P}_n be the set of univariate polynomials of degree n and $\bar{\mathcal{P}}_n \stackrel{\text{def}}{=} \mathcal{P}_0 \cup \dots \cup \mathcal{P}_n$ the set of univariate polynomials of degree at most n . For instance, \mathcal{P}_0 is the set of constant functions, \mathcal{P}_1 the set of linear functions, and $\bar{\mathcal{P}}_1 = \mathcal{P}_0 \cup \mathcal{P}_1$ the set of constant and linear functions. We have:

$$\delta \in \mathcal{P}_n \Rightarrow \{\delta^{(e)} = 0 \mid e > n\} \quad (1)$$

$$\delta^{(e)} = 0 \Rightarrow \{\delta \in \mathcal{P}_n \mid n < e\} \quad (2)$$

$$\delta^{(e)} = 0 \Leftrightarrow \delta \in \bar{\mathcal{P}}_{e-1}. \quad (3)$$

3.3 Univariate Rational Functions

A univariate rational function μ is a fraction of two polynomials of degrees $a, b \in \mathbb{N}$ respectively. We define μ 's degree as the pair $(a, b) \in \mathbb{N}^2$. Formally, a function μ is a univariate rational function if and only if

$\exists a, b \in \mathbb{N}$ s.t. $\mathcal{R}_{a,b}[\mu]$ with:

$$\mathcal{R}_{a,b}[\mu] \stackrel{\text{def}}{\Leftrightarrow} \left(\exists \alpha \in \mathcal{P}_a, \gamma \in \mathcal{P}_b, \text{gcd}(\alpha, \gamma) \in \mathbb{R}^* \text{ s.t. } \mu = \frac{\alpha}{\gamma} \right), \quad (4)$$

where $\text{gcd}(\alpha, \gamma)$ is the polynomial greatest common divisor of α and γ . The condition on $\text{gcd}(\alpha, \gamma)$ means that the fraction cannot be simplified to a lower degree. The domain of μ is $\Omega \subset \mathbb{R}$ with $\Omega = \{x \in \mathbb{R} \mid \gamma(x) \neq 0\}$.

A function μ is a univariate rational function of degree at most (a, b) if $\bar{\mathcal{R}}_{a,b}[\mu]$ holds with:

$$\bar{\mathcal{R}}_{a,b}[\mu] \stackrel{\text{def}}{\Leftrightarrow} \left(\exists \alpha \in \bar{\mathcal{P}}_a, \gamma \in \bar{\mathcal{P}}_b \text{ s.t. } \mu = \frac{\alpha}{\gamma} \right). \quad (5)$$

Abusing notation, we write $\mu \in \mathcal{R}_{a,b}$ and $\mu \in \bar{\mathcal{R}}_{a,b}$ equivalently to $\mathcal{R}_{a,b}[\mu]$ and $\bar{\mathcal{R}}_{a,b}[\mu]$ respectively. The set of univariate rational functions forms a field, meaning that the addition, subtraction, multiplication and division of rational functions is a rational function.

4 Canonical Invariants of Univariate Rational Functions

We show that a univariate rational function μ of degree (a, b) satisfies differential invariants. These invariants form an infinite set for a given degree (a, b) , as function μ also satisfies the invariants of higher degrees. We here define and study the *canonical invariant* of degree (a, b) . Importantly, when applied to a function $\mu \in C^\infty$, the canonical invariant constrains μ to be a univariate rational function of degree at most (a, b) . It is straightforward to see that the canonical invariant thus implies all the invariants of higher degrees.

Our key result is given in the next proposition. It is an equivalence between the set of univariate rational functions of some maximal degree and the canonical differential invariant of that degree. The proof of this proposition is given as appendix B.

Proposition 1 (Canonical invariants). *A function μ is a univariate rational function of degree at most (a, b) if and only if it satisfies the canonical invariant $\mathcal{I}_{a,b}$. Formally:*

$$\forall \mu \in C^\infty, a, b \in \mathbb{N} \quad \bar{\mathcal{R}}_{a,b}[\mu] \Leftrightarrow \mathcal{I}_{a,b}[\mu]$$

with:

$$\mathcal{I}_{a,b}[\mu] \stackrel{\text{def}}{\Leftrightarrow} \left(\sum_{s \in \mathcal{S}_{b+1}} \text{sgn}(s) \prod_{i=1}^{b+1} \binom{a+i}{s_i-1} \mu^{(a+i-s_i+1)} = 0 \right). \quad (6)$$

We use \mathcal{S}_{b+1} to generate all permutations of the set $[1, b+1]$ and $\text{sgn}(s) \in \{-1, 1\}$ to compute the signature¹

¹We have that $\text{sgn}(s) = 1$ if s can be generated by an even number of interchanges of the elements of $[1, b+1]$, and $\text{sgn}(s) = -1$ otherwise.

of permutation $s = \{s_1, \dots, s_{b+1}\}$. $\mathcal{I}_{a,b}[\mu]$ is an homogeneous polynomial ODE of degree $b + 1$. The lowest and greatest orders it involves are $\max(a - b + 1, 0)$ and $a + b + 1$ respectively.

Examples. We give four examples of univariate rational functions and canonical invariants, and discuss the general case and advanced examples. An important case is the one of linear fractional functions, which are univariate rational functions of degree $(1, 1)$. Any univariate rational function μ of degree $(1, 1)$ can be written as $\mu(x) = \frac{a_1x+a_0}{b_1x+b_0}$ for $a_0, a_1, b_0, b_1 \in \mathbb{R}$. Proposition 1 says that such functions are also characterized by the following differential canonical invariant:

$$\mathcal{I}_{1,1}[\mu] \Leftrightarrow \left(-2\mu^{(1)}\mu^{(3)} + 3\mu^{(2)}\mu^{(2)} = 0 \right).$$

This can be easily verified. This invariant is the ODE from which the Schwarzian derivative was derived (Ovsienko and Tabachnikov, 2005), and as expected, it corresponds to linear fractional functions, called homographies in the context of projective geometry. The same reasoning leads to the differential canonical invariant for the quadratic fractional functions, using $(a, b) = (2, 2)$, as:

$$\begin{aligned} \mathcal{I}_{2,2}[\mu] \Leftrightarrow & \left(-12\mu^{(1)}\mu^{(3)}\mu^{(5)} + 15\mu^{(1)}\mu^{(4)}\mu^{(4)} \right. \\ & + 18\mu^{(2)}\mu^{(2)}\mu^{(5)} - 30\mu^{(2)}\mu^{(4)}\mu^{(3)} \\ & \left. + 40\mu^{(3)}\mu^{(3)}\mu^{(3)} - 30\mu^{(3)}\mu^{(2)}\mu^{(4)} = 0 \right). \end{aligned}$$

The numerator's and denominator's degrees need not be equal, and we can for instance derive the canonical invariants for the constant-quadratic or the quadratic-linear fractional functions as:

$$\begin{aligned} \mathcal{I}_{0,2}[\mu] \Leftrightarrow & \left(\mu^{(0)}\mu^{(0)}\mu^{(3)} - 3\mu^{(0)}\mu^{(2)}\mu^{(1)} \right. \\ & \left. + 6\mu^{(1)}\mu^{(1)}\mu^{(1)} - 3\mu^{(1)}\mu^{(0)}\mu^{(2)} = 0 \right), \end{aligned}$$

and:

$$\mathcal{I}_{2,1}[\mu] \Leftrightarrow \left(-3\mu^{(2)}\mu^{(4)} + 4\mu^{(3)}\mu^{(3)} = 0 \right).$$

We observe that the order of the invariants increases with a , while the number of terms and degrees increase with b .

Generally speaking, any rational plane curve induces a rational warp which thus satisfies the proposed invariants. For example, the warp μ induced by a cubic B-spline satisfies $\mathcal{I}_{3,3}[\mu]$, while the warp μ induced by a cubic NURBS satisfies $\mathcal{I}_{9,9}[\mu]$, as shown in §5.3.

Polynomial kernel of the canonical invariants. The kernel of the simple derivative invariant of order e is formed by the polynomials of degree $e - 1$, see equation (3). The canonical invariant of order (a, b) also includes polynomials in its kernel, forming the polynomial kernel of the canonical invariant. Their degree is now derived in terms of a, b . As can be seen from equation (6) and from the examples above, the canonical invariant $\mathcal{I}_{a,b}[\mu]$ is a sum where each term has partial derivatives of μ as factors. Therefore, using equation (1), there must exist a polynomial $\hat{\mu}$ so that $\mathcal{I}_{a,b}[\hat{\mu}]$ holds. We call $\hat{\mu} \in \bar{\mathcal{P}}_d$ the *polynomial kernel* of $\mathcal{I}_{a,b}$, and define the *polynomial kernel's degree* d as the greatest natural number so that $\forall \hat{\mu} \in \bar{\mathcal{P}}_d, \mathcal{I}_{a,b}[\hat{\mu}]$ holds. It is easy to show that $d = a$. This is obtained by considering that each term in $\mathcal{I}_{a,b}[\mu]$ must vanish. Because a term is a product of derivatives, what matters is thus the greatest order involved in the term, as the derivative at this order vanishes with a higher degree polynomial than the other derivatives. Overall, for all terms to vanish, we must thus select $d + 1$ as the minimum over the terms in $\mathcal{I}_{a,b}[\mu]$ of the greatest order involved in each term. By examining equation (6) we can compute the sum of orders on each term as $\sum_{i=1}^{b+1} (a + i - s_i + 1) = (a + 1)(b + 1)$, which is a constant value. Additionally, there is a term in $\mathcal{I}_{a,b}[\mu]$ where all orders are equal. Together, these imply that $d + 1$, the minimum over the terms of the greatest order per term, is given by the order of the term where all factors have the same order. Because we have $b + 1$ factors per term, we obtain $d = \frac{(a+1)(b+1)}{b+1} - 1 = a$. For the four examples above we have $d = 1, d = 2, d = 0$ and $d = 2$. In the first example, it says that the polynomial kernel is made of linear functions $\hat{\mu} \in \mathcal{P}_1$, for which $\hat{\mu}^{(2)} = 0$. Of course the canonical invariant does not only have polynomials in its kernel, but also univariate rational functions, according to proposition 1. The polynomial kernel is indeed included in the general kernel, as $\hat{\mu} \in \bar{\mathcal{P}}_a \Rightarrow \bar{\mathcal{R}}_{a,0}[\hat{\mu}] \Rightarrow \bar{\mathcal{R}}_{a,b}[\hat{\mu}]$ for $b \geq 0$.

5 Deformable Single-View Geometry

We first give our notation and show how the warp is derived for a generic embedding and perspective projection. We then specialize the derivation to rational embeddings, showing that the warp is then itself rational. We finally show how to estimate the warp under a locally-rational regularization constructed from the canonical invariants.

5.1 Notation

We denote the deformed plane curve as \mathcal{C} in the projective plane \mathbb{P}^2 . We represent \mathcal{C} by a projective embedding $\varphi : \Omega \rightarrow \mathbb{P}^2$ where $\Omega \subset \mathbb{R}$ is the curve parameterization's domain, leading to $\mathcal{C} = \varphi(\Omega)$. In other words, we have that \mathcal{C} is the set of points of the projective plane $\mathbf{Q} = \varphi(p) \in \mathbb{P}^2$ obtained by embedding any point $p \in \Omega$. The camera is given by a function $\Pi : \mathbb{P}^2 \rightarrow \mathbb{R}$ representing a 1D projective camera (Quan and

Kanade, 1997), and the warp is simply $\eta \stackrel{\text{def}}{=} \Pi \circ \varphi$. We require $\eta \in C^\infty(\Omega, \mathbb{R})$, meaning that the warp must be a smooth function, except at some points where the invariants may not be evaluated. This is implied by requiring that \mathcal{C} is a smooth curve lying entirely in the open projective half-plane representing the front or the back of the camera, but however forms a weaker requirement.

5.2 The Warp for a Generic Embedding

The general 1D projective camera Π can be defined as the composition of a 2D homography $h : \mathbb{P}^2 \rightarrow \mathbb{P}^2$ and the canonical perspective projection as:

$$\Pi(\mathbf{Q}) \stackrel{\text{def}}{=} \frac{\mathbf{h}_1^\top \mathbf{Q}}{\mathbf{h}_2^\top \mathbf{Q}}, \quad (7)$$

with $\mathbf{h}_1, \mathbf{h}_2 \in \mathbb{R}^3$ and not both identically zero. These 6 parameters are defined up to scale, meaning that only 5 of them are independent, and contain the camera's 2 intrinsics and 3 extrinsics. Whether the camera is calibrated or not does not change the way deformable SVG is defined in this context. Using equation (7) we derive the warp as:

$$\eta = \Pi \circ \varphi = \frac{\mathbf{h}_1^\top \varphi}{\mathbf{h}_2^\top \varphi}. \quad (8)$$

5.3 The Warp for a Locally-Rational Embedding

We denote the three components of the projective embedding as $\varphi^\top = [\varphi_1 \ \varphi_2 \ \varphi_3]$. We first assume that the embedding is rational, meaning that we have $\varphi_i \in \bar{\mathcal{R}}_{a_i, b_i}$, and so that there exist $\alpha_i \in \bar{\mathcal{P}}_{a_i}$, $\gamma_i \in \bar{\mathcal{P}}_{b_i}$, $i \in \{1, 2, 3\}$ with $\varphi_i = \frac{\alpha_i}{\gamma_i}$. We analyse the numerator and denominator of η in equation (8) by letting $l \in \{1, 2\}$ as:

$$\begin{aligned} \mathbf{h}_l^\top \varphi &= h_{l1}\varphi_1 + h_{l2}\varphi_2 + h_{l3}\varphi_3 \\ &= h_{l1} \frac{\alpha_1}{\gamma_1} + h_{l2} \frac{\alpha_2}{\gamma_2} + h_{l3} \frac{\alpha_3}{\gamma_3} \\ &= \frac{h_{l1}\alpha_1\gamma_2\gamma_3 + h_{l2}\alpha_2\gamma_3\gamma_1 + h_{l3}\alpha_3\gamma_1\gamma_2}{\gamma_1\gamma_2\gamma_3}. \end{aligned}$$

We thus have:

$$\frac{\mathbf{h}_1^\top \varphi}{\mathbf{h}_2^\top \varphi} = \frac{h_{11}\alpha_1\gamma_2\gamma_3 + h_{12}\alpha_2\gamma_3\gamma_1 + h_{13}\alpha_3\gamma_1\gamma_2}{h_{21}\alpha_1\gamma_2\gamma_3 + h_{22}\alpha_2\gamma_3\gamma_1 + h_{23}\alpha_3\gamma_1\gamma_2}.$$

Since $\delta_1 \in \bar{\mathcal{P}}_a, \delta_2 \in \bar{\mathcal{P}}_b \Rightarrow \delta_1\delta_2 \in \bar{\mathcal{P}}_{a+b}, \delta_1 + \delta_2 \in \bar{\mathcal{P}}_{\max(a,b)}$ we have $\eta \in \bar{\mathcal{R}}_{a,b}$ with:

$$a = b = \max(a_1 + b_2 + b_3, a_2 + b_3 + b_1, a_3 + b_1 + b_2),$$

where we recall that (a_i, b_i) is the degree of φ_i , $i \in \{1, 2, 3\}$. Therefore, we have $\eta \in \bar{\mathcal{R}}_{a,a}$, and we can use proposition 1 to state that $\mathcal{I}_{a,a}[\eta]$ must hold. This way, we have obtained constraints that the warp must satisfy in order to be geometrically valid for a rational embedding. This brings an answer to point (ii) of SVG, which we defined in the introduction as a study of the warp’s characteristics, independently of the embedding and projection. Our target assumption however is that the embedding is locally-rational, meaning that it is rational in an infinitesimal neighbourhood of each point $p \in \Omega$. Consequently, the warp is also locally-rational. The key advantage is that representing a non-trivial curve with a globally rational embedding requires one to use high-degree polynomials for its numerator, its denominator or both of them, whereas using high-degree polynomials is not stable in practice (Press et al., 2007). In contrast, a locally-rational warp is a flexible and stable warp, which will be constrained to behave similarly to a low-degree rational warp on a local basis. A similar reasoning may be found in the definition of B-spline curves, which are both flexible and stable thanks to the use of local low-degree polynomials (Piegl and Tiller, 1997). In practice, we use a stable generic deformable model for the warp and perform warp estimation using a regularizer constructed from the canonical invariants. When used as a penalty term in warp estimation, this regularizer makes the warp behave locally like a rational map. This process is described in the next section, where we use a simple gaussian radial basis function to model the warp.

5.4 Estimating the Warp from Correspondences

We describe how the canonical invariants fit into a framework to estimate the warp from point correspondences. This framework preserves the generic warp model’s natural flexibility while making it locally-rational. It is different from using an explicit rational embedding which would make the warp globally rational.

General methodology. We want to estimate the warp η from m point correspondences $\{p_j, q_j\}$, $p_j \in \Omega$, $q_j \in \mathbb{R}$, $j \in [1, m]$. The general methodology is to define a correspondence cost functional $C[\eta] = \frac{1}{m} \sum_{j=1}^m (\eta(p_j) - q_j)^2$ and a regularizer $R[\eta]$, and to solve the following variational problem:

$$\min_{\eta} C[\eta] + \lambda R[\eta], \quad (9)$$

where $\lambda \in \mathbb{R}^+$ is the regularization weight, controlling the relative influence of the two terms in relationship to the intrinsic scale of each term.

Regularization weight. The value of λ is important. We choose it by splitting the correspondences into a training, a validation and a test set. We then sample λ on a predefined range, solve problem (9) for each

value of λ using the training set and finally keep the warp with the lowest validation residual. The sampling range is chosen such that outside it the resulting warp does not exhibit significant changes. We use 30 samples in our experiments.

Regularizer. In the literature, the regularizer is often derived from the L_2 norm of the warp’s partial derivatives integrated over the domain. For example, the B-spline (Piegl and Tiller, 1997) and the Thin-Plate Spline (Duchon, 1976) were constructed by penalizing second-order derivatives. We call these the polynomial regularizers because they constrain the warp to be locally-polynomial. We use a low order $k \in \{1, 2, 3\}$ because the regularizers are differential, thus measured very locally. This already allows significant flexibility on the warp. They are defined as:

$$R_{\text{POL},k}[\eta] \stackrel{\text{def}}{=} \int_{\Omega} \left(\eta^{(k)}(p) \right)^2 dp.$$

We however show in §5.3 that a locally-polynomial or rational embedding leads to a locally-rational warp with degree (a, a) . We thus use the canonical rational invariant $\mathcal{I}_{a,a}$ given in proposition 1 to construct new regularizers, for $a \in \{1, 2, 3\}$, as:

$$R_{\text{RAT},a}[\eta] \stackrel{\text{def}}{=} \int_{\Omega} (\mathcal{I}_{a,a}[\eta](p))^2 dp.$$

The regularizers are applied to the whole warp. However, they are constructed as integrals over the domain Ω . Importantly, the integrands measure the warp’s closeness to a polynomial or a rational function at point $p \in \Omega$. The regularizers thus measure the extent to which the warp is everywhere locally similar to a polynomial or a rational function. In other words, the lower the regularizer, the closer the warp to a local polynomial or rational function on each and every point of its domain.

Minimization. In order to solve problem (9), we introduce a flexible parametric representation of the warp using a gaussian radial basis function (Bishop, 2006). The l gaussian bases are uniformly spread across the domain Ω , which we define as $\Omega = \{p \in \mathbb{R} \mid 0 \leq p \leq 1\}$ by default. Their position are defined as $c_1, \dots, c_l \in \Omega$ and thus given by $c_k = \frac{k-1}{l-1}$, $k \in [1, l]$. Their standard deviation σ is fixed. The parameters to estimate are thus contained in a set $\omega \in \mathbb{R}^l$ of l weights, one for each gaussian basis. We use $l = 50$ gaussian bases with $\sigma = \frac{5}{l}$ in our experiments. The parametric warp representation is thus $\eta(p; \omega) \stackrel{\text{def}}{=} \sum_{k=1}^l \omega_k \exp\left(-\frac{(p-c_k)^2}{2\sigma^2}\right)$. We use gradient descent starting from the identity warp.

Results. Each estimated warp is tied to a regularizer and thus denoted accordingly, as POL1, POL2, POL3, RAT1, RAT2 and RAT3. As a baseline, we also used the radial basis function representation fitted without a regularizer, by setting $\lambda = 0$ in problem (9). The resulting warp is denoted PLAIN.

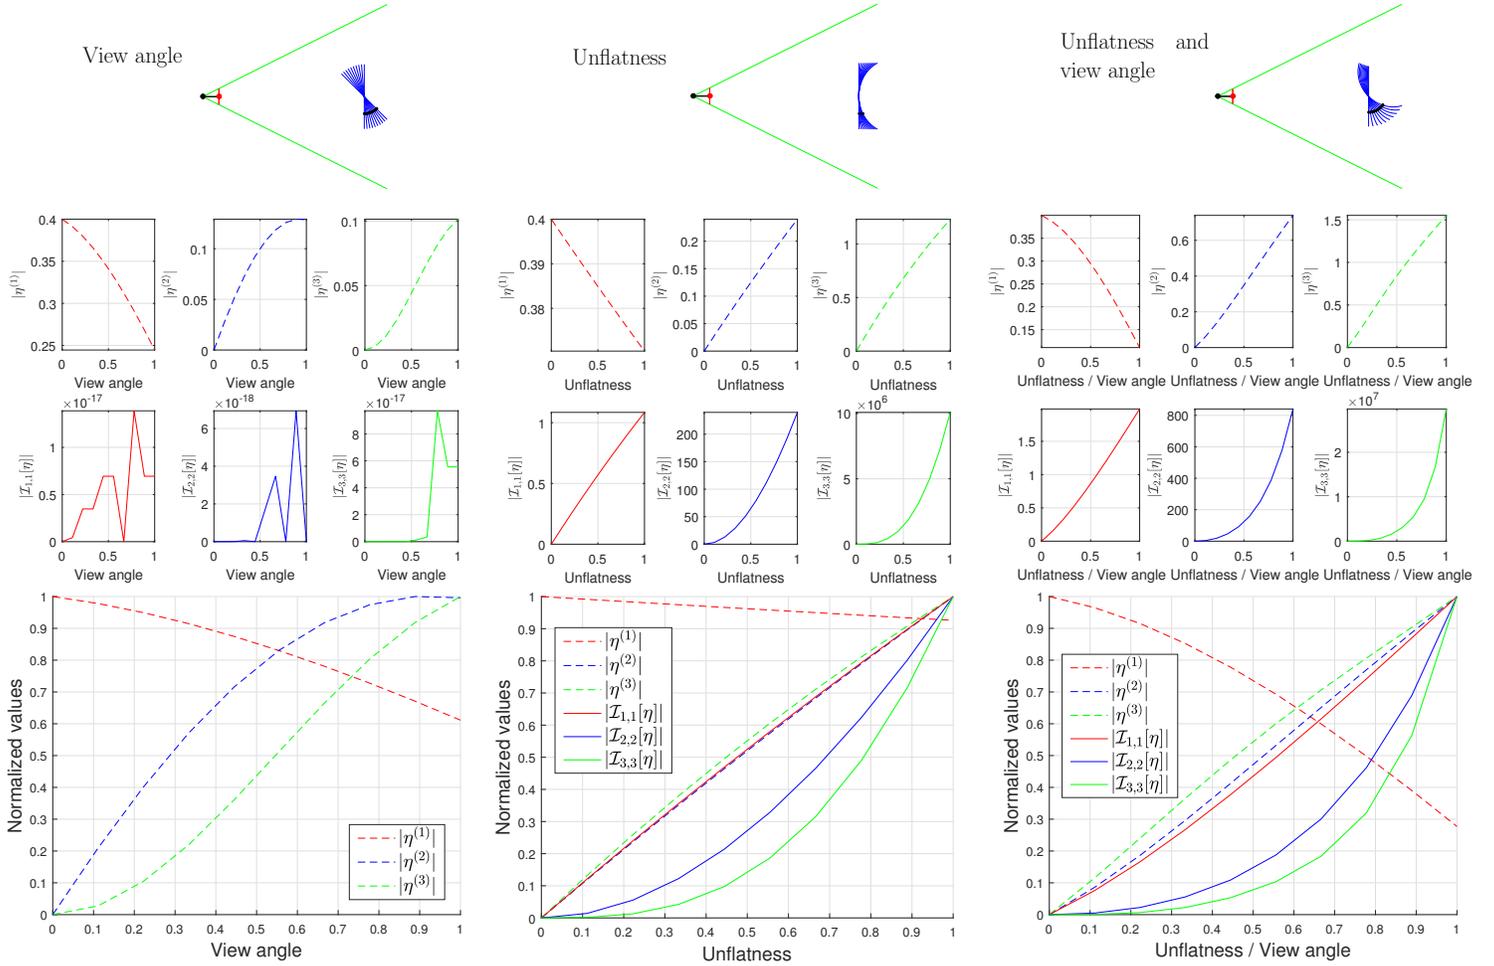


Figure 2: Results of simulated data experiments. From left to right, the object is slanted, curved and both slanted and curved. From top to bottom, the simulated camera and object with the observed point shown in black, individual plots and joint plot with normalization of the six measured quantities.

6 Experimental Results

Our two main goals are *(i)* to assess to which extent the invariants hold on an independent warp and *(ii)* to assess their contribution to the stability of warp estimation when used to form regularizers. We report two sets of experiments to evaluate these two main goals, both using simulated and real data.

6.1 Assessing the Invariants on Independent Warps

We want to compare the value of the invariants estimated for independent warps, in order to objectively find the ability of each invariant to capture the warp’s behavior. This is because, even if theoretically the rational invariants use a more physically valid representation of the warp than the simple warp derivatives, their dependency on higher-order derivatives may cause instabilities. For simulated data, we use the true warp, while for real data, we use the PLAIN warp. We first describe the measured quantities.

6.1.1 Measured Quantities

We measured six quantities related to the six regularizers defined in §5.4. These are the first three derivatives of the warp, namely $|\eta^{(1)}|$, $|\eta^{(2)}|$ and $|\eta^{(3)}|$, related to the classical polynomial regularizers, and the first three canonical invariants, namely $|\mathcal{I}_{1,1}[\eta]|$, $|\mathcal{I}_{2,2}[\eta]|$ and $|\mathcal{I}_{3,3}[\eta]|$, related to the rational regularizers. We use the absolute value as only the magnitude of these quantities matters, not their sign. However, because they are of different orders of magnitude, they cannot be cross-compared individually. We thus also report a normalized version of these six quantities achieved by dividing by the maximum over a set of values.

6.1.2 Simulated Data

Simulation setup. We simulated a 1D perspective camera with a 1 unit focal length observing a 1D flat object of size 4 units whose midpoint is located 10 units from the centre of projection. The flat object is then deformed, which is represented by simulating its embedding from the template domain Ω . By projecting the object, we arrive at a simulated image. Because the warp η is the composition of the known projection and embedding functions, as shown by equation (8), we can estimate the warp’s target value and its derivatives analytically at any point in Ω . We use a fixed point shown in black on the simulated shape in figure 2 to monitor the six measured quantities. We then did three types of basic transformations of the object: slanting, curving and a combination of both. The embedding is rational for the former but not for the last two transformations, as it involves trigonometric functions. This way, our measurements depend directly on the basic transformations of slanting and curving, and can be plotted according to the amount of perspective, of curvature, and of their combination.

Perspective. The default fronto-parallel view of the object does not contain perspective. In order to increase the perspective effect, we slanted the object by rotating it around its midpoint with an angle θ varying between 0 and 45 degrees. The results are shown in the left column of figure 2. We observe that the warp derivatives are sensitive to perspective. More precisely, $|\eta^{(2)}|$ and $|\eta^{(3)}|$ vanish at $\theta = 0$, when the object is exactly fronto-parallel, but then quickly increase with perspective, while $|\eta^{(1)}|$ does not even vanish at $\theta = 0$. On the other hand, all three canonical invariants $|\mathcal{I}_{1,1}[\eta]|$, $|\mathcal{I}_{2,2}[\eta]|$ and $|\mathcal{I}_{3,3}[\eta]|$ vanish independently of the amount of perspective.

Curvature. The default flat-shaped object does not contain curvature. In order to increase curvature, we curved the object by linearly morphing its shape with an arc of radius 4 units. Importantly, the resulting shape is not described by a polynomial or a rational function. The results are shown in the middle column of figure 2. We make the same observation on the warp derivatives as in the increasing perspective case: these are sensitive to curvature. On the other hand, contrarily to the case of perspective, the canonical invariants also show sensitivity to curvature. We observe that the sensitivity decreases with the order. More precisely, $|\mathcal{I}_{1,1}[\eta]|$ performs very similarly to $|\eta^{(2)}|$ and $|\eta^{(3)}|$, while $|\mathcal{I}_{2,2}[\eta]|$ and $|\mathcal{I}_{3,3}[\eta]|$ show better performance, $|\mathcal{I}_{3,3}[\eta]|$ having a flat regime approximately twice as long as $|\mathcal{I}_{2,2}[\eta]|$. We observe that $|\mathcal{I}_{2,2}[\eta]|$ and $|\mathcal{I}_{3,3}[\eta]|$ have a large intrinsic scale, exhibiting larger values than the other four invariants. The intrinsic scale of an invariant is arbitrary. Importantly however, it does not influence the results of warp estimation, thanks to the cross-validation mechanism selecting the regularization weight automatically.

Combined perspective and curvature. We increased both perspective and curvature by combining slanting and curving of the object as above-described. The results are shown in the right column of figure 2. We make the same observations as in the increasing curvature case: all six measured quantities are sensitive to the combined effect of perspective and curvature. There are two important differences however. First, there is a larger gap between the group formed by the warp derivatives and $|\mathcal{I}_{1,1}[\eta]|$ and the group formed by $|\mathcal{I}_{2,2}[\eta]|$ and $|\mathcal{I}_{3,3}[\eta]|$. Second, there is a smaller gap between $|\mathcal{I}_{2,2}[\eta]|$ and $|\mathcal{I}_{3,3}[\eta]|$.

Synthesis. Our experimental observations perfectly confirm the theory. First, we have that the warp derivatives are not invariant to camera perspective and object curvature. Second, we have that all canonical invariants are insensitive to camera perspective and have various sensitivities to object curvature. The first-order canonical invariant is very sensitive to object curvature while the second- and third-order canonical invariants have a good tolerance. This is sensible because a higher order canonical invariant models a higher order rational embedding able to represent the object’s local curvature.

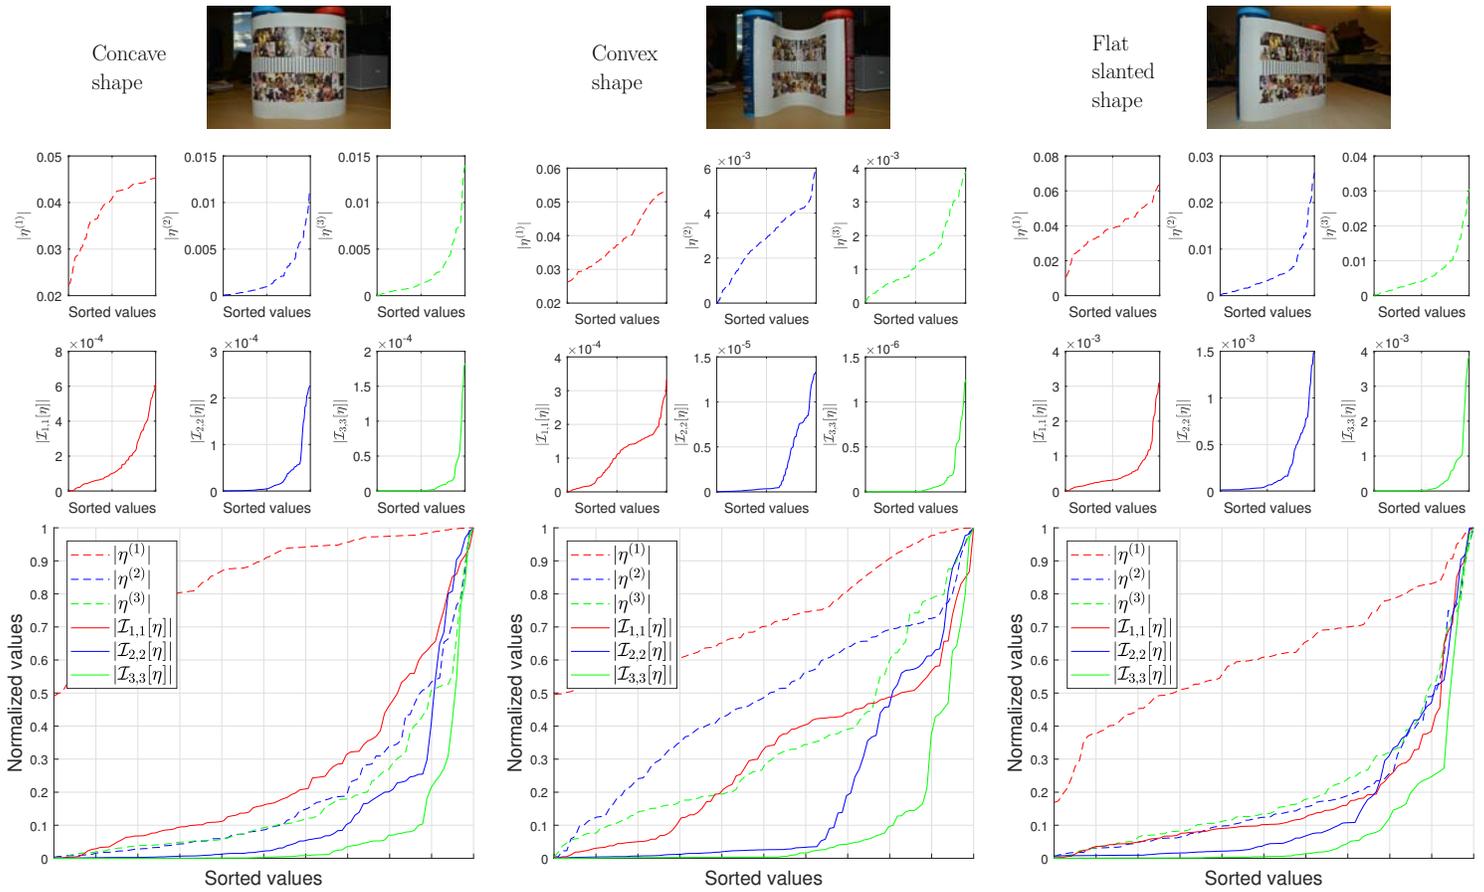


Figure 3: Results of real data experiments. The top row shows the input 2D images from which the 1D images were constructed. The middle and bottom rows show the six measured quantities individually and in a joint plot with normalization, respectively.

6.1.3 Real Data

Data acquisition and result presentation. We emulated a 1D camera by slicing space along a plane containing the centre of projection of a 2D camera. Concretely, we created a 1D template with 30 regularly spaced black tick marks which we printed on a paper sheet. We then selected a horizontal line crossing the ticks on the 2D image of the deformed paper sheet. The ticks were manually marked in the 2D image and these locations projected on the line. The 1D coordinates of the ticks along this line define the 1D image obeying the geometry of a 1D camera. Each tick thus gives a 1D correspondence between the template and the image. The interest of this dataset lies in the noise distribution on the data, which is more realistic than on simulated data, in the realism of the 2D shape and in the realism of the physical imaging process, to which the pin-hole camera is just an approximation. We fitted the warp PLAIN to the correspondences. We then chose 100 points uniformly spread in the template and used the warp to estimate their location in the image and their derivatives. From these we estimated the value of the six measured quantities, sorted them, and plotted them individually and jointly with normalization, as for the simulated data. For a particular quantity, the plotted curve thus starts with the lowest values, obtained for the points where the warp best minimizes the quantity, and increases towards the highest values, obtained for the points where the warp is least compatible with the quantity.

Results. The results obtained for three images are shown in figure 3. The first-order derivative $|\eta^{(1)}|$ does clearly not do a good job, as it has a steady logarithmic or linear increase. The second- and third-order derivatives $|\eta^{(2)}|$ and $|\eta^{(3)}|$ perform better, but also show an increase at all points. Similarly, the first-order canonical invariant $|\mathcal{I}_{1,1}[\eta]|$ show an increase at all points, performing worse than $|\eta^{(2)}|$ and $|\eta^{(3)}|$. The second- and third-order canonical invariants $|\mathcal{I}_{2,2}[\eta]|$ and $|\mathcal{I}_{3,3}[\eta]|$ however both have a flat part starting at the beginning of their curves. This flat part means that these two quantities, and thus the corresponding invariants, form a better model of the warp than the other quantities. The third-order canonical invariant $|\mathcal{I}_{3,3}[\eta]|$ is significantly better than the second-order one $|\mathcal{I}_{2,2}[\eta]|$. These results agree with the synthetic data experiments to some extent. Indeed, as predicted by theory, we have that the canonical invariants perform better than the mere derivatives, because they model the effect of perspective. The results also differ from the synthetic data experiments as for the canonical invariants, the larger the degree the better the results. This is explained by the fact that the local shape is probably of higher complexity in the real data case, and thus requires a more complex warp to be accurately captured.

6.2 Assessing the Invariants in Warp Estimation

We want to evaluate the performance of the six regularizers defined in §5.4 in warp estimation from point correspondences.

6.2.1 Simulated Data

We describe the simulation setup and then the results obtained when varying the amount of noise, the number of points and when creating an extrapolation area.

Simulation setup. We used the same simulation setup as in §6.1.2 with a more complex shape obtained by combining sines and cosines with a global rotation and translation. An example of shape is shown in figure 4. We generated a finite number of point correspondences which we corrupted with noise following a gaussian distribution whose standard deviation controls the noise level. We used a default noise level of 0.5% of the object’s image size and number of points of 20. We splitted the correspondence set equally in a training and validation sets. We used the true warp to assess the estimated warps by measuring the discrepancy at 1,000 regularly sampled test points at the zeroth-, first- and second-orders. The zeroth-order discrepancy is the distance between the position predicted by the true and the estimated warp. The first- and second-order discrepancies are the differences between the first- and second-order derivatives predicted by the true and the estimated warp. The results are averages over 50 trials.

Noise. We vary the noise level between 0 to 1% of the object’s image size. The results are shown in the left column of figure 4. We observe that the overall accuracy does not show the typical trend of degrading with noise. This is explained by the fact that regularization smoothes out noise. We observe that the POL1 warp gives the poorest performance. It is followed by a group comprising the PLAIN, POL2 and POL3 warps, with POL2 performing slightly better in terms of first-order error. We next find the RAT2 and RAT3 warps followed by the RAT1 warp. These three warps are on par in terms of zeroth-order error but the RAT1 warp performs clearly best in terms of the first- and second-order errors.

Number of points. We vary the number of points between 10 and 50. The results are shown in the middle column of figure 4. We observe that the accuracy increases with the number of points for all warps at zeroth- and first-order. However, this increase becomes very shallow for the POL1, POL2, POL3 and PLAIN warps at second-order. The POL1 warp has the worst performance, followed by a group with the POL2, POL3 and PLAIN warps. The RAT2 and RAT3 warps perform better and the RAT1 warp outperforms, except for lower numbers of points. The differences between the warps increase with the order.

Extrapolation. We created an extrapolation domain by removing input correspondences in a localized

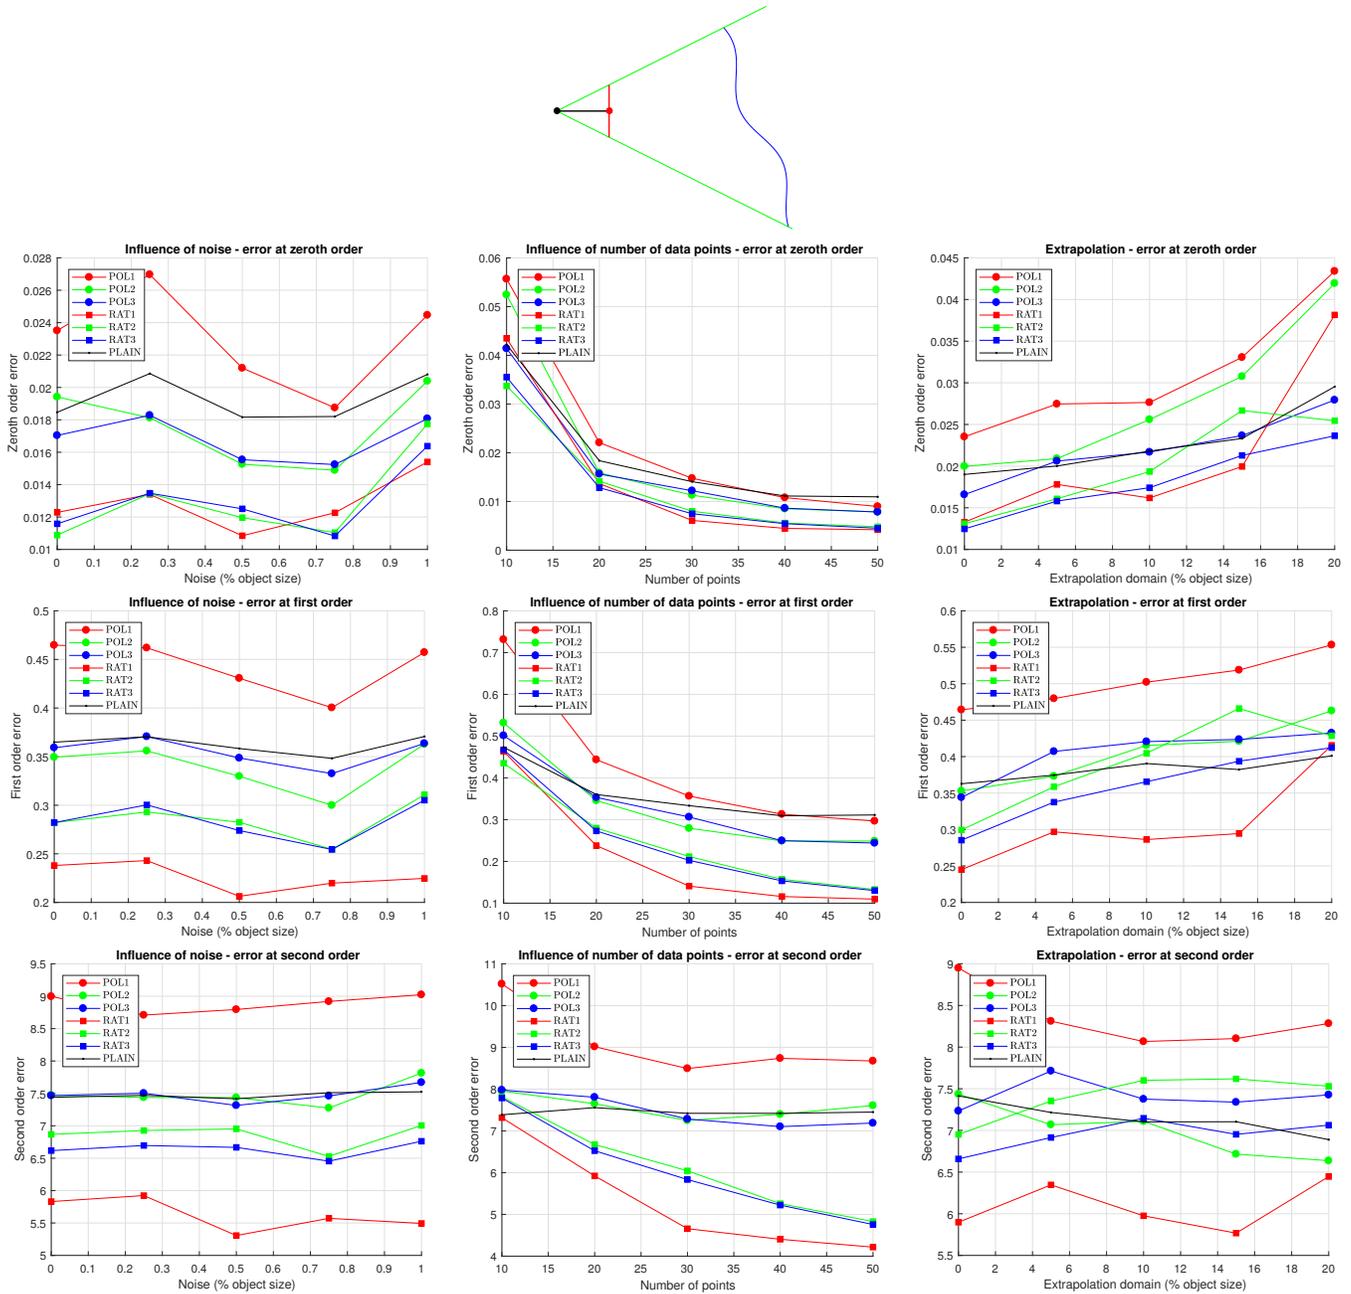


Figure 4: Results of simulated data experiments. The top row shows an example of simulated configuration. From left to right, the noise, number of points and extrapolation domain are varied. From the second to fourth row, the zeroth-, first- and second-order errors are shown for the seven tested warps.

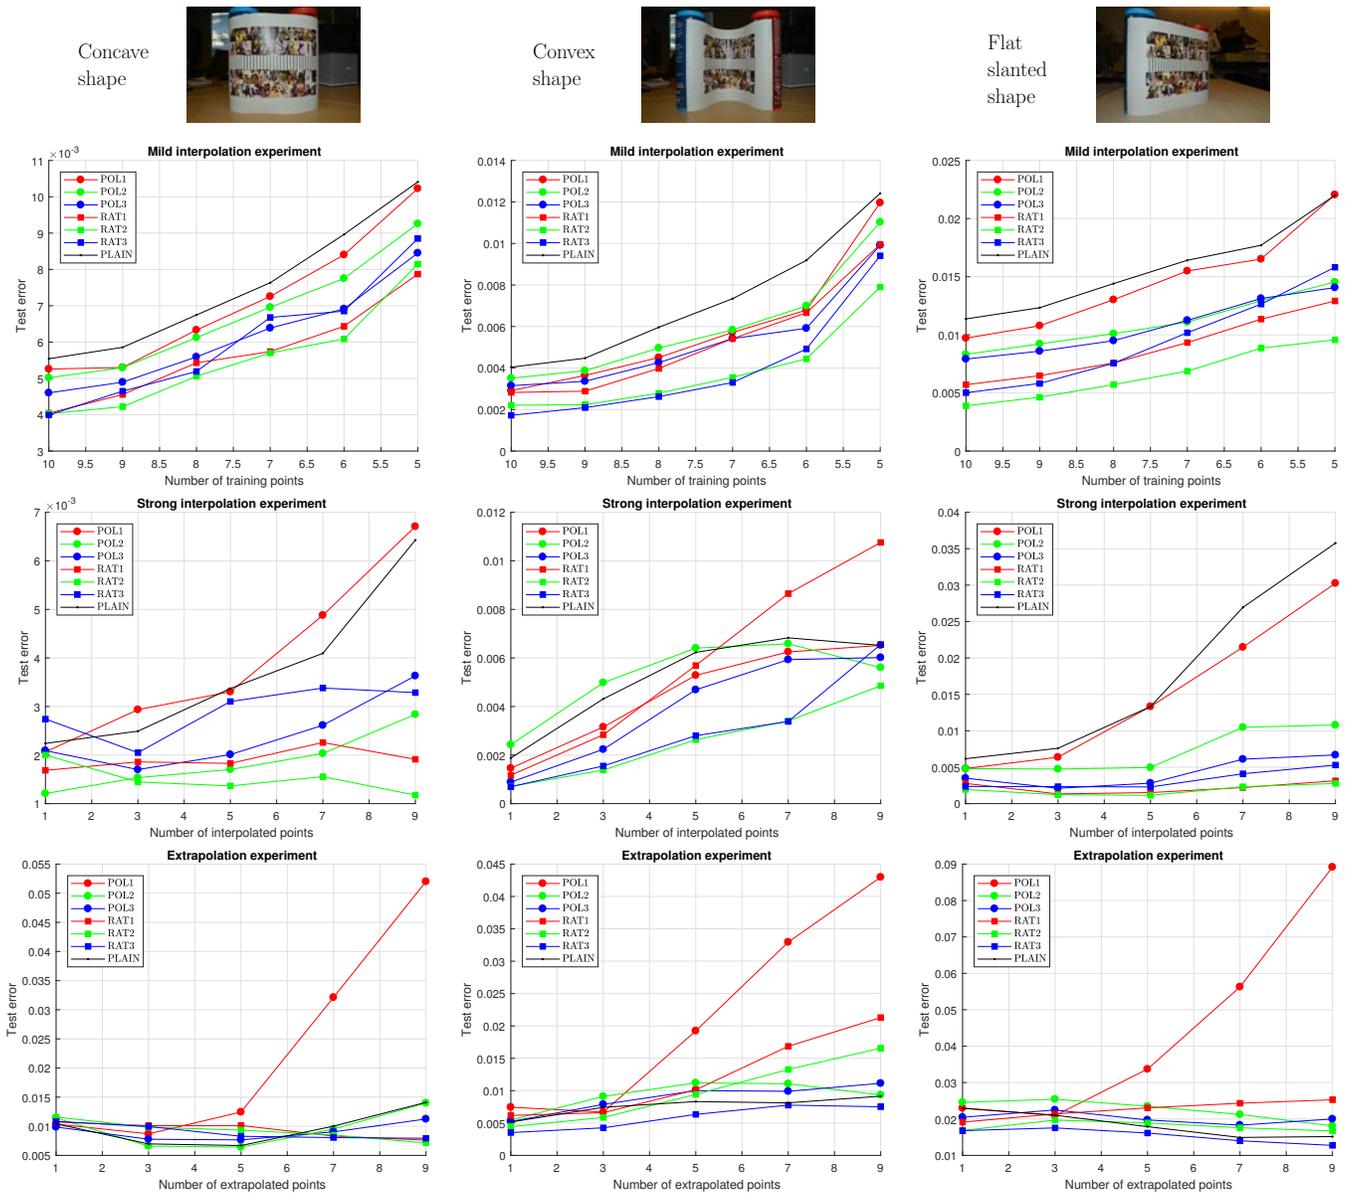


Figure 5: Results of real data experiments. The top row shows the input 2D images from which the 1D images were constructed. The other three rows show the result of the mild interpolation, strong interpolation and extrapolation experiments, respectively.

part on the side of the object. We observe that the accuracy degrades with the size of the extrapolation domain. The POL1 warp has the worst performance and the RAT1 warp has the best performance. We find inbetween the five other warps with performances depending on the order and the size of the extrapolation domain. In particular, the RAT3 warp always outperforms at zeroth- and first-orders. It outperforms at second-order when the extrapolation domain is smaller than 10% of the object.

Synthesis. Our experimental observations show that there is a real benefit in using a rational regularizer over a regular polynomial regularizer in warp estimation. We have that the first-order rational regularizer outperforms in almost all cases. It is closely followed by the second- and third-order rational regularizers, then by the second- and third-order polynomial regularizers. Interestingly, the plain gaussian radial basis function is on par with the second- and third-order rational regularizers. Finally, the first-order polynomial regularizer underperforms in almost all cases. The differences across the warps generally increase with the order of the estimated derivatives. The reason for these results is that the first-order rational regularizer forms the optimal trade-off in terms of local shape capture capability and local constraint strength. Indeed, the higher the degree, the more flexible the modeling capability but the weaker the constraint. The powerful modeling capability of the second- and third-order rational regularizers is thus impaired by the weakened local constraint they exert on the warp, compared to the first-order rational regularizer.

6.2.2 Real Data

In the case of real data, we do not have access to ground truth. We followed the estimation methodology described in §5.4 and three cases of splitting the 30 input point correspondences into training, validation and test sets. This allowed us to emulate mild and strong interpolation and extrapolation. The results are the test residuals given in figure 5. They are averages over 50 trials for each configuration in each case.

Mild interpolation. The mild interpolation case emulates the non-uniformity of data by changing the size of the training set, without creating overly large areas of missing data. We used a fixed size of 10 points for each of the validation and test sets while vaying the number of training points from 10 to 5. These points are all chosen randomly at each trial. The results are shown in the second row of figure 5. We observe that all warps follow the same trend of nicely degrading for a decreasing number of training points. The PLAIN warp always performs worst while the RAT2 warp always performs best, except for the convex shape where the RAT3 warp performs slightly better for larger numbers of training points. The second worst is the POL1 warp. It is followed by the POL2 and then the POL3 warps. The POL3 warp has performances close to the RAT1 warp, which is itself below the RAT2 warp. Overall, the RAT2 warp has the best performance. This is because the RAT2 warp has increased modeling capabilities compared to the RAT1 warp from its higher

degree. However, because the interpolation range is limited, the extra modeling capabilities of the RAT3 warp did not compensate its lower constraining strength on the warp.

Strong interpolation. The strong interpolation case emulates the lack of data in possibly large areas surrounded by available data by changing the size of the test set within the domain. We split the input points in three parts. The left and right parts are used to sample the training and validation sets while the middle part forms the test set. We use a fixed number of 10 points for each of the training and validation sets and ensure that 5 of each lies in the left and right parts. The number of test points is then varied from 1 to 9. It represents the size of the gap that the warp has to interpolate. The results are shown in the third row of figure 5. We observe that all warps roughly follow the same trend of degrading for an increasing interpolation range. The PLAIN and POL1 warps are the worst, except for the convex shape where the POL2 and RAT1 warps also perform poorly. The POL3 and RAT3 warps show better performances. The RAT1 warp is on par, except for the convex shape. The RAT2 warp performs better in all cases, except a single configuration with one point on the concave shape. Overall, the RAT2 warp has the best performance. The explanation for these results are similar to the mild interpolation case. Even though the interpolation range is larger, this remains a relatively easy task for the warps, and the extra modeling capability of the RAT3 warp is not needed.

Extrapolation. The extrapolation case emulates the lack of data on the boundary of the domain by changing the size of the test set. We vary the number of test points from 1 to 9 and choose them as the left-most points in the input points. We use a fixed number of 10 points for each of the training and validation sets. They are chosen randomly in the remaining input points. The results are shown in the fourth row of figure 5. We observe a mild error increase or a slight error decrease, depending on the shape, for all warps but the POL1 warp, for an increasing extrapolation range. Indeed, the POL1 warp quickly degrades dramatically in all three shapes. It is then difficult to give a general ordering of the warps, as their performance changes significantly across the configurations and shapes. Nonetheless, we observe that the PLAIN warp does well as compared to the two previous cases of interpolation. While the best results is generally achieved by the RAT3 warp, the POL2 and POL3 warps outperform on the concave shape for smaller extrapolations. Overall, the RAT3 warp has the best performance. This is explained by the fact that the RAT3 warp has the highest modeling capabilities of all the tested warps, and that extrapolation is fundamentally a more difficult task than interpolation, in the sense that the model must capture the structure of the data to a very good extent in order to make reliable predictions.

7 Conclusion

We have proposed a theoretical framework based on locally-rational embeddings and univariate rational functions to formulate the SVG of a deformable body observed by a 1D projective camera. This framework is generic in the sense that it does not impose a specific deformation constraint, but uses mere local-smoothness represented by a locally-rational embedding. Our work is theoretic and includes experimental results showing the ability of the proposed invariants to improve the estimation of warps. Our framework fits in the recent and fascinating research topic of understanding the visual geometry of a deformable body. We now discuss two possibilities of future work based on our framework: deformable MVG and higher dimensions.

A natural question when working on deformable SVG is whether the framework would form a basis for deformable MVG. Local smoothness surely forms a good basis for deformable MVG because locally-smooth warps have a local group structure. Concretely for two images, deformable SVG studies the smooth warps $\eta_i \in C^\infty$ from the parameterization space defined from the embeddings as $\eta_i = \Pi \circ \varphi_i$, $i = 1, 2$, while deformable MVG studies the smooth inter-image warp $\eta_{1,2} \in C^\infty$. By construction we have $\eta_{1,2} = \eta_2 \circ \eta_1^{-1} = \Pi \circ \varphi_2 \circ \eta_1^{-1}$. We can thus construct a smooth embedding $\varphi_{1,2}$ by using image 1 as a parameterization space for image 2 as $\varphi_{1,2} \stackrel{\text{def}}{=} \varphi_2 \circ \eta_1^{-1}$. Therefore, one could use local-smoothness to define the warps and exploit the proposed deformable SVG.

Extending our 1D framework to higher dimensions also forms an appealing idea. The extension to 2D is particularly interesting, as it will allow one to handle the 2D images of surfaces taken by regular cameras.

Acknowledgments

This research has received funding from the EU's FP7 through the ERC research grant 307483 FLEXABLE. We thank the authors of (Gallardo et al., 2015) for the real dataset and Yan Gérard for his kind feedback on the paper.

References

- A. Bartoli, Y. Gérard, F. Chadebecq, T. Collins, and D. Pizarro. Shape-from-template. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(10):2099–2118, October 2015.
- C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer-Verlag, 2006.

- F. L. Bookstein. Principal warps: Thin-Plate Splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, June 1989.
- J. Duchon. Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *RAIRO Analyse Numérique*, 10:5–12, 1976.
- R. Fabbri and B. B. Kimia. Multiview differential geometry of curves. *International Journal of Computer Vision*, 117(3):1–23, December 2016.
- O. Faugeras. *Three-Dimensional Computer Vision*. MIT Press, 1993.
- O. Faugeras and B. Mourrain. On the geometry and algebra of the point and line correspondences between n images. In *International Conference on Computer Vision*, 1995.
- O. Faugeras, L. Quan, and P. Sturm. Self-calibration of a 1d projective camera and its application to the self-calibration of a 2d projective camera. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(10):1179–1185, October 2000.
- D. Forsyth and J. Ponce. *Computer Vision – A Modern Approach*. Pearson, 2012. International Edition.
- D. Fortun, P. Bouthemy, and C. Kervrann. Optical flow modeling and computation: A survey. *Computer Vision and Image Understanding*, 134:1–21, May 2015.
- M. Gallardo, D. Pizarro, A. Bartoli, and T. Collins. Shape-from-template in flatland. In *International Conference on Computer Vision and Pattern Recognition*, 2015.
- N. A. Gumerov, A. Zandifar, R. Duraiswami, and L. S. Davis. 3D structure recovery and unwarping surfaces applicable to planes. *International Journal of Computer Vision*, 66(3):261–281, 2006.
- R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2003. Second Edition.
- J. Kaminski and A. Shashua. Multiple view geometry of general algebraic curves. *International Journal of Computer Vision*, 56(3):195–219, 2004.
- V. Ovsienko and S. Tabachnikov. *Projective Differential Geometry Old and New*. Cambridge University Press, 2005.
- V. Ovsienko and S. Tabachnikov. What is... the schwarzian derivative? *Notices of the AMS*, 56(1), January 2009.

- S. Parashar, D. Pizarro, and A. Bartoli. Isometric non-rigid shape-from-motion with riemannian geometry solved in linear time. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(10):2442–2454, October 2018.
- M. Perriollat, R. Hartley, and A. Bartoli. Monocular template-based reconstruction of inextensible surfaces. *International Journal of Computer Vision*, 95(2):124–137, November 2011.
- L. Piegl and W. Tiller. *The NURBS Book*. Monographs in Visual Communication. Springer-Verlag, 1997. Second Edition.
- D. Pizarro, R. Khan, and A. Bartoli. Schwarzs: Locally projective image warps based on 2D schwarzian derivatives. *International Journal of Computer Vision*, 119(2):93–109, September 2016.
- W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes – The Art of Scientific Computing*. Cambridge University Press, 2007. Third Edition.
- L. Quan and T. Kanade. Affine structure from line correspondences with uncalibrated affine cameras. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(8):834–845, August 1997.
- M. Salzmann, J. Pilet, S. Ilic, and P. Fua. Surface deformation models for nonrigid 3D shape recovery. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(8):1–7, August 2007.
- C. Schmid and A. Zisserman. The geometry and matching of lines and curves over multiple views. *International Journal of Computer Vision*, 40(3):199–234, 2000.
- J. R. Sendra, F. Winkler, and S. Prez-Daz. *Rational Algebraic Curves A Computer Algebra Approach*. Springer, 2008.
- R. J. Walker. *Algebraic Curves*. Springer-Verlag, 1978.

A Higher Dimensions

We give general points about extending our framework to higher dimensions. We then show that the 1D invariants provide a partial set of invariants for the 2D setup using a virtual 1D setup.

A.1 General Points

Extending our 1D framework to higher dimensions present different types of difficulties, whether one extends the dimension of the source or the target space. Note that extending the framework to handle 2D warps

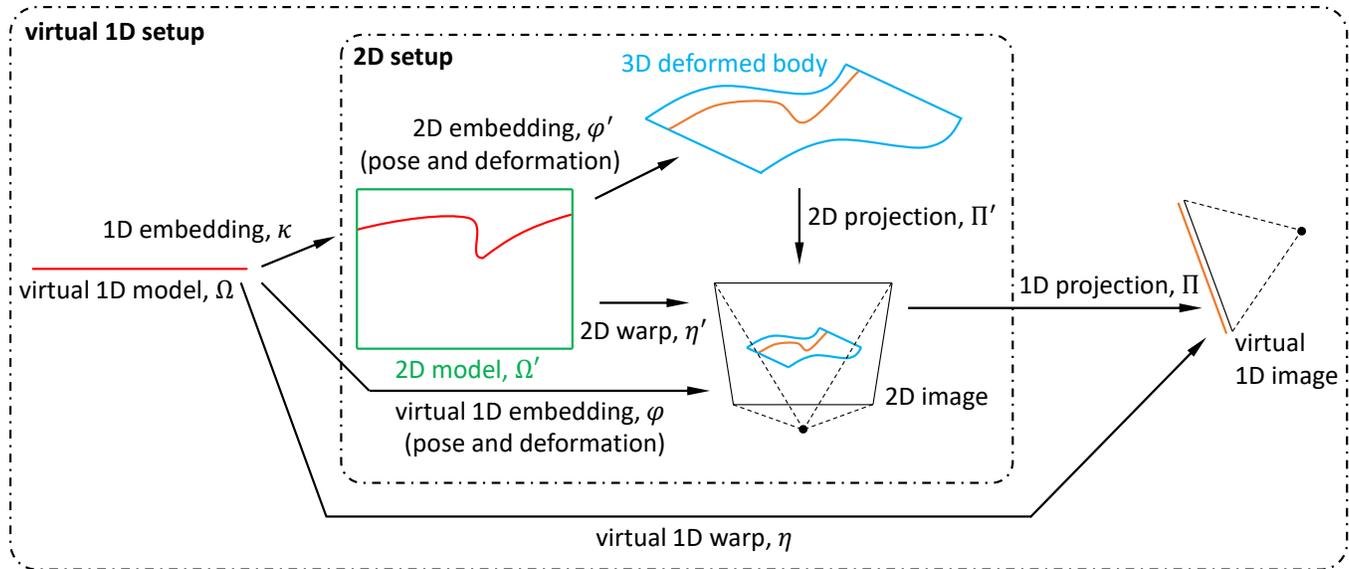


Figure 6: The virtual 1D setup. We use this construction to show that the proposed invariants for the 1D setup form a subset of the invariants for the 2D setup.

requires one to extend the dimension of both the source and target spaces. Extending the dimension of the target space, for example to 2D in order to model the image of a deformable 3D curve, leads to $\eta^\top = [\eta_1 \ \eta_2] = \frac{1}{\gamma}[\alpha_1 \ \alpha_2]$. It is straightforward to see that this warp satisfies two invariants, $\mathcal{I}_{a,a}[\eta_1]$ and $\mathcal{I}_{a,a}[\eta_2]$, for a polynomial embedding of degree a . However, these two invariants, according to proposition 1, guarantee that $\eta_i = \frac{\alpha_i}{\gamma_i}$ for some $\alpha_i, \gamma_i \in \mathcal{P}_a$, $i = 1, 2$, but they do not guarantee that $\gamma_1 = \gamma_2$, which is a necessary condition for the SVG to be valid. It is quite likely that this requirement adds an invariant that η must fulfill. Extending the dimension of the source space, for example to 2D in order to model a deformable 3D surface, means that the partial derivatives will become tensor-valued functions. Though this will require one to use a tensorial notation, such as Einstein's, this will also allow one to follow the proposed mathematical framework.

A.2 The Virtual 1D Setup

Our goal is to show that the 1D invariants form a subset of the 2D invariants. For that purpose we construct a virtual 1D setup, as shown in figure 6. We start from the 2D setup, shown in figure 1, where the warp $\eta' : \Omega' \rightarrow \mathbb{R}^2$, $\Omega' \subset \mathbb{R}^2$, is a 2D rational mapping. Therefore, both its source and target space dimensions must be reduced by one in order to apply our 1D invariants. Importantly, these reductions must preserve rationality. Our first step is to reduce the source space dimension. This is done by choosing a rational plane curve parameterized by $\kappa : \Omega \rightarrow \mathbb{R}^2$, a 1D rational embedding. Clearly, $\eta' \circ \kappa : \Omega \rightarrow \mathbb{R}^2$ is a rational function. The domain $\Omega \subset \mathbb{R}$ of κ stands for the virtual 1D body's model. Our second step is to reduce the target

space dimension. This is done by introducing a virtual 1D camera, represented by $\Pi : \mathbb{P}^2 \rightarrow \mathbb{R}$, in the 2D image. Clearly, $\eta = \Pi \circ \eta' \circ \kappa : \Omega \rightarrow \mathbb{R}$ is a rational function. Therefore, according to proposition 1, there exist $a, b \in \mathbb{N}$ such that the invariant $\mathcal{I}_{a,b}[\eta]$ holds. The value of a, b is upper-bounded by a combination of the degrees of κ and η' and the parameters of Π . We can easily derive interesting cases out of this general formulation. For instance, choosing the rational plane curve as a straight horizontal line in the 2D model, we have $\kappa(p) = [p \quad y]^\top$, where $y \in \mathbb{R}$ is a constant, and the degree of $\eta' \circ \kappa$ is at most (a', b') , the degree of η' . Then, placing the virtual 1D camera at the origin with its principal axis along the vertical direction, we have $\Pi([q_1 \quad q_2]^\top) = \frac{q_1}{q_2}$, and the degree of η becomes at most (a', a') . The real example shown in the experiments of §§6.1.3 and 6.2.2 represents a special case of this general principle.

We have shown that the 1D invariants can be applied to simple restrictions of the 2D warp, based on choosing a simple rational curve in the 2D model and a simple 1D projection in the 2D image. This forms a possible way to study the generalization of our 1D invariants to higher dimensions. A first step would be to study how to form a minimal set of restrictions amongst the many possible ones to form a set of mutually independent invariants. A second step would be to study if the so-obtained invariant set forms sufficient conditions to characterize the set of 2D rational functions.

B Proof of Proposition 1

We prove proposition 1 in two parts. Part I is the forward implication and part II the reverse implication. Part I requires the following three lemmas. It constructs the canonical invariant $\mathcal{I}_{a,b}[\mu]$ by differentiating the relationship $\gamma\mu = \alpha$, called $\mathcal{G}[\mu, \alpha, \gamma]$ in lemma 1, at the orders $\mathcal{Z} \subset \mathbb{N}$, which we will show can be chosen as $\mathcal{Z} = [a + 1, a + b + 1]$ in lemma 2. It then combines the different orders in lemma 3 to eliminate all orders of γ and α from the equations, resulting in the sought invariant, an equation depending on μ only.

Lemma 1. *We have $\forall \mu \in C^\infty, a, b \in \mathbb{N}, \alpha \in \bar{\mathcal{P}}_a, \gamma \in \bar{\mathcal{P}}_b$:*

$$\mu = \frac{\alpha}{\gamma} \Rightarrow \{\mathcal{H}_{e,b}[\mu, \gamma] \mid e > a\},$$

with:

$$\mathcal{H}_{e,b}[\mu, \gamma] \stackrel{\text{def}}{\Leftrightarrow} \left(\sum_{k=0}^b \binom{e}{k} \gamma^{(k)} \mu^{(e-k)} = 0 \right). \quad (10)$$

Proof. We have that $\mu = \frac{\alpha}{\gamma}$ implies $\mathcal{G}[\mu, \alpha, \gamma]$, with:

$$\mathcal{G}[\mu, \alpha, \gamma] \stackrel{\text{def}}{\Leftrightarrow} (\gamma\mu = \alpha). \quad (11)$$

By taking the e -th derivative of $\mathcal{G}[\mu, \alpha, \gamma]$ using Leibniz's rule we obtain:

$$D^e \mathcal{G}[\mu, \alpha, \gamma] \Leftrightarrow \left(\sum_{k=0}^e \binom{e}{k} \gamma^{(k)} \mu^{(e-k)} = \alpha^{(e)} \right). \quad (12)$$

Using equation (1), we have that $\alpha^{(e)} = 0$ for $e > a$ and $\gamma^{(e)} = 0$ for $e > b$. We can thus rewrite $D^e \mathcal{G}[\mu, \alpha, \gamma]$ as:

$$\mathcal{K}_{e,a,b}[\mu, \alpha, \gamma] \stackrel{\text{def}}{\Leftrightarrow} \left(\sum_{k=0}^{\min(e,b)} \binom{e}{k} \gamma^{(k)} \mu^{(e-k)} = \mathbb{1}_{e \leq a} \alpha^{(e)} \right),$$

where $\mathbb{1}$ is the indicator function, with $\mathbb{1}_{\text{true}} = 1$ and $\mathbb{1}_{\text{false}} = 0$. For $e > a$ we can simplify $\mathcal{K}_{e,a,b}[\mu, \alpha, \gamma]$ to $\mathcal{H}_{e,b}[\mu, \gamma]$, using the relationship $\min(e, b) \leq b$ and $\binom{e}{k} = 0$ for $k > e$ to set a fixed summation count. \square

Lemma 2. For $\alpha \in \bar{\mathcal{P}}_a, \gamma \in \bar{\mathcal{P}}_b$ constructing a differential invariant $\mathcal{I}_{a,b}[\mu]$ requires one to differentiate $\mathcal{G}[\mu, \alpha, \gamma]$ at the orders in $\mathcal{Z} \subset \mathbb{N}$ with $\min(\mathcal{Z}) > a$ and $|\mathcal{Z}| = b + 1$.

Proof. Recall that deriving a differential invariant $\mathcal{I}_{a,b}[\mu]$ on μ from equations $D^e \mathcal{G}[\mu, \alpha, \gamma]$ formed at several orders e requires that α and γ be eliminated, as well as their derivatives at all orders. We start with the elimination of α and its derivatives. By construction $D^e \mathcal{G}[\mu, \alpha, \gamma]$ involves $\alpha^{(e)}$ and no other orders of α , as can be seen in equation (12). Therefore, the only way to cancel out all orders of α to form an invariant on μ is by differentiation of $\mathcal{G}[\mu, \alpha, \gamma]$ at an order greater than a , the degree of α , thanks to equation (1). This proves $\min(\mathcal{Z}) > a$. Importantly, once α has been cancelled, the equations $D^e \mathcal{G}[\mu, \alpha, \gamma]$ become homogeneous, as can trivially be seen from equation (12). We now turn to the elimination of γ and its derivatives. The situation is different as each equation $D^e \mathcal{G}[\mu, \alpha, \gamma]$ depends on $\{\gamma^{(0)}, \dots, \gamma^{(e)}\}$. This means that we will have to form a linear system in γ and its derivatives. Forcing this system to have a solution will then allow us to eliminate γ and its derivatives, as will be shown in the proof of lemma 3. This also means that differentiating $D^e \mathcal{G}[\mu, \alpha, \gamma]$ to form $D^{e+1} \mathcal{G}[\mu, \alpha, \gamma]$ introduces a new order $\gamma^{(e+1)}$ of γ . Therefore the first e orders in $\mathcal{G}[\mu, \alpha, \gamma]$ involve the first e orders of γ . This means that to eliminate all orders of γ to form an invariant on μ one has to differentiate $\mathcal{G}[\mu, \alpha, \gamma]$ to an order greater than b , the degree of γ , thanks to equation (1), and that the number of orders involved must be greater than b . This is equivalent to requiring the size of \mathcal{Z} to be greater than b . Each order beyond b then allows one to form a linear system, and thus an invariant. Because we want to form one invariant only, this proves $|\mathcal{Z}| = b + 1$. \square

Lemma 3. Defining $\hat{\mathcal{H}}_{a,b}[\mu, \gamma] \stackrel{\text{def}}{=} \{\mathcal{H}_{e,b}[\mu, \gamma] \mid e \in [a + 1, a + b + 1]\}$, we have:

$$\forall \mu, \gamma \in C^\infty, a, b \in \mathbb{N}^* \quad \hat{\mathcal{H}}_{a,b}[\mu, \gamma] \Rightarrow \mathcal{I}_{a,b}[\mu].$$

Proof. We choose $\mathcal{U}, \mathcal{V} \subset \mathbb{N}$ to contain respectively the orders of μ and γ involved in equation (10), namely $\mathcal{U} = [\max(a+1-b, 0), a+b+1]$ and $\mathcal{V} = [0, b]$. The relationship $\mathcal{H}_{e,b}[\mu, \gamma]$ in equation (10) is bilinear in $\mu^{(\mathcal{U})}$ and $\gamma^{(\mathcal{V})}$ and we thus have $\forall e \in \mathbb{N}, a, b \in \mathbb{N}, \exists \mathbf{A}_{e,b} \in \mathbb{N}^{|\mathcal{U}| \times |\mathcal{V}|}$:

$$\mathcal{H}_{e,b}[\mu, \gamma] \Leftrightarrow \left(\mu^{(\mathcal{U})\top} \mathbf{A}_{e,b} \gamma^{(\mathcal{V})} = 0 \right).$$

Together with $\hat{\mathcal{H}}_{a,b}[\mu, \gamma]$ this shows that $\gamma^{(\mathcal{V})}$ is up to scale a non trivial element in the kernel of the matrix-valued function $\beta_{a,b}[\mu] \in C^\infty(\mathbb{R}, \mathbb{R}^{(b+1) \times |\mathcal{V}|})$ whose ‘rows’ are functions $\mu^{(\mathcal{U})\top} \mathbf{A}_{e,b} \in C^\infty(\mathbb{R}, \mathbb{R}^{1 \times |\mathcal{V}|})$ for $e \in [a+1, a+b+1]$. Because $|\mathcal{V}| = b+1$, we thus have $\forall \mu, \gamma \in C^\infty, a, b \in \mathbb{N}, \exists \beta_{a,b}[\mu] \in C^\infty(\mathbb{R}, \mathbb{R}^{(b+1) \times (b+1)})$:

$$\hat{\mathcal{H}}_{a,b}[\mu, \gamma] \Leftrightarrow \left(\beta_{a,b}[\mu] \gamma^{([0,b])} = 0 \right).$$

In order to form a non-trivial invariant on μ , we force the kernel of $\beta_{a,b}[\mu]$ to be non-empty. This invariant may thus be found by expanding $\det(\beta_{a,b}[\mu]) = 0$. Inspecting $\mathcal{H}_{e,b}[\mu, \gamma]$ in equation (10), we find that the entry of matrix $\beta_{a,b}[\mu]$ at (i, j) is given by:

$$\beta_{a,b;i,j}[\mu] = \binom{a+i}{j-1} \mu^{(a+i-j+1)}. \quad (13)$$

Using Leibniz’s formula we finally expand $\det(\beta_{a,b}[\mu]) = 0$ to arrive at the canonical invariant (6). \square

Proof of proposition 1, part I. We want to show the forward implication, $\bar{\mathcal{R}}_{a,b}[\mu] \Rightarrow \mathcal{I}_{a,b}[\mu]$. We use the definition (4) to obtain $\bar{\mathcal{R}}_{a,b}[\mu] \Rightarrow \exists \alpha \in \bar{\mathcal{P}}_a, \gamma \in \bar{\mathcal{P}}_b$ s.t. $\mu = \frac{\alpha}{\gamma}$. From lemma 1, we then have that this implies $\{\mathcal{H}_{e,b}[\mu, \gamma] \mid e > a\}$. Lemma 2 shows that instantiating $\mathcal{H}_{e,b}$ at the orders in $[a+1, a+b+1]$, which is equivalent to $\hat{\mathcal{H}}_{a,b}[\mu, \gamma]$, is required to form the canonical invariant $\mathcal{I}_{a,b}[\mu]$. These orders lead to the canonical invariant and any other orders respecting lemma 2 lead to a non-canonical invariant for the degree (a, b) (see comment 1). Finally, applying lemma 3 proves the forward implication in equation (6). \square

Comment 1 (Non-canonical invariants and number of invariants). *The way the invariants are constructed depends on $\mathcal{Z} \subset \mathbb{N}$, which must satisfy the constraints from lemma 2. We have defined the canonical invariant as the one obtained with $\mathcal{Z} = [a+1, a+b+1]$. However, for a given degree (a, b) , there exist an infinite number of \mathcal{Z} subsets, each leading to a different invariant. There thus exist an infinite number of invariants.*

Part II of our proof of proposition 1 uses lemma 4 to show that given the invariant $\mathcal{I}_{a,b}[\mu]$ there exists a smooth vector-value function $\gamma \in (C^\infty)^{b+1}$, $\gamma^\top \stackrel{\text{def}}{=} (\gamma_1, \dots, \gamma_{b+1})$ with $\gamma_k \in C^\infty$, $k = 1, \dots, b+1$, so that $\hat{\mathcal{H}}_{a,b}[\mu, \gamma]$ holds. The latter is defined as $\hat{\mathcal{H}}_{a,b}[\mu, \gamma] \stackrel{\text{def}}{=} \{\mathcal{H}_{e,b}[\mu, \gamma] \mid e \in [a+1, a+b+1]\}$, where $\mathcal{H}_{e,b}[\mu, \gamma]$

is an extension of $\mathcal{H}_{e,b}[\mu, \gamma]$ where γ 's derivatives are replaced by the scalar-valued functions forming the vector-valued function γ as:

$$\mathcal{H}_{e,c}[\mu, \gamma] \Leftrightarrow \left(\sum_{k=0}^c \binom{e}{k} \gamma_{k+1} \mu^{(e-k)} = 0 \right).$$

It then shows that this implies that at each point in Ω either μ or γ is a polynomial of some bounded degree. In the former case μ is also a univariate rational function. In the latter case, this is because the entries of γ correspond to the derivatives of a scalar-valued function γ . We can then show that there exist a polynomial α such that $\mu = \frac{\alpha}{\gamma}$. More specifically, our proof requires the following three lemmas.

Lemma 4. *We have $\forall \mu \in C^\infty, a, b \in \mathbb{N}$:*

$$\mathcal{I}_{a,b}[\mu] \Rightarrow \left(\exists \gamma \in (C^\infty)^{b+1} \text{ s.t. } \hat{\mathcal{H}}_{a,b}[\mu, \gamma] \right) \tag{14}$$

$$\mathcal{I}_{a,b}[\mu] \Rightarrow \left(\exists \gamma \in C^\infty \text{ s.t. } \hat{\mathcal{H}}_{a,b}[\mu, \gamma^{([0,b])}] \right). \tag{15}$$

Proof. We first prove equation (14). The way the invariant $\mathcal{I}_{a,b}$ was derived in the proof of lemma 3 is by nullifying the determinant of a matrix-valued operator $\beta_{a,b}$ on μ . In other words, $\mathcal{I}_{a,b}[\mu]$ implies $\det(\beta_{a,b}[\mu]) = 0$, and thus that $\beta_{a,b}[\mu]$ has a non-empty null-space, which we represent by the vector-valued function γ , defined such that $\beta_{a,b}[\mu]\gamma = 0$. The smoothness of γ is implied by the smoothness of μ and $\beta_{a,b}[\mu]$, and the fact that the null-space of a rank-deficient matrix is obtained as a multi-linear combination of the matrix' entries.

We now prove equation (15) by contradiction. Assume that $\forall \mu \in C^\infty$ s.t. $\mathcal{I}_{a,b}[\mu]$ we have $\nexists \gamma \in C^\infty$ s.t. $\hat{\mathcal{H}}_{a,b}[\mu, \gamma^{([0,b])}]$. By using the contrapositive of lemma 1, this implies that $\nexists \alpha \in \mathcal{P}_a, \gamma \in \mathcal{P}_b$ s.t. $\mu = \frac{\alpha}{\gamma}$. In other words, this implies the falsity of $\mathcal{R}_{a,b}[\mu]$. However, the forward implication of proposition 1, which is proved in part I of the proof, implies that the kernel of the canonical invariant comprises at least the rational functions of the appropriate degree, which contradicts the falsity of $\mathcal{R}_{a,b}[\mu]$. \square

Lemma 5. *We have $\forall \mu, \gamma \in C^\infty, c, e \in \mathbb{N}, e \geq c$:*

$$(\mathcal{H}_{e,c}[\mu, \gamma] \wedge \mathcal{H}_{e+1,c}[\mu, \gamma]) \Rightarrow \gamma^{(c+1)} \mu^{(e-c)} = 0.$$

Proof. We first differentiate $\mathcal{H}_{e,c}[\mu, \gamma]$:

$$D\mathcal{H}_{e,c}[\mu, \gamma] \Leftrightarrow \left(\sum_{k=0}^c \binom{e}{k} \gamma^{(k)} \mu^{(e-k+1)} + \sum_{k=0}^c \binom{e}{k} \gamma^{(k+1)} \mu^{(e-k)} = 0 \right).$$

We expand the first term as:

$$\begin{aligned} & \sum_{k=0}^c \binom{e}{k} \gamma^{(k)} \mu^{(e-k+1)} \\ &= \binom{e}{0} \gamma \mu^{(e+1)} + \sum_{k=1}^c \binom{e}{k} \gamma^{(k)} \mu^{(e-k+1)} \\ &= \binom{e+1}{0} \gamma \mu^{(e+1)} + \sum_{k=1}^c \binom{e}{k} \gamma^{(k)} \mu^{(e-k+1)}, \end{aligned}$$

where the last equality follows from $\binom{e}{0} = \binom{e+1}{0} = 1$ since $e \geq 0$. We expand the second term as:

$$\begin{aligned} & \sum_{k=0}^c \binom{e}{k} \gamma^{(k+1)} \mu^{(e-k)} \\ &= \sum_{k=1}^{c+1} \binom{e}{k-1} \gamma^{(k)} \mu^{(e-k+1)} \\ &= \sum_{k=1}^c \binom{e}{k-1} \gamma^{(k)} \mu^{(e-k+1)} + \binom{e}{c} \gamma^{(c+1)} \mu^{(e-c)}. \end{aligned}$$

Using the relation $\binom{e+1}{k} = \binom{e}{k-1} + \binom{e}{k}$ we obtain:

$$D\mathcal{H}_{e,c}[\mu, \gamma] \Leftrightarrow \left(\sum_{k=0}^c \binom{e+1}{k} \gamma^{(k)} \mu^{(e-k+1)} + \binom{e}{c} \gamma^{(c+1)} \mu^{(e-c)} = 0 \right).$$

Expanding $D\mathcal{H}_{e,c}[\mu, \gamma] - \mathcal{H}_{e+1,c}[\mu, \gamma]$ yields $\binom{e}{c} \gamma^{(c+1)} \mu^{(e-c)} = 0$. Because $e \geq c$, $\binom{e}{c} > 0$ and we arrive at $\gamma^{(c+1)} \mu^{(e-c)} = 0$. \square

Lemma 6. We have $\forall \mu \in C^\infty, c, e \in \mathbb{N}, \gamma \in C^\infty, \gamma^{(c+1)} = 0$:

$$\mathcal{H}_{e,c}[\mu, \gamma] \Rightarrow (\gamma \mu)^{(e)} = 0.$$

Proof. Let $\alpha = \gamma \mu$. We thus have that $\mathcal{G}[\mu, \alpha, \gamma]$ from equation (11) holds. We then form $D^e \mathcal{G}[\mu, \alpha, \gamma]$ from

equation (12) and $D^e \mathcal{G}[\mu, \alpha, \gamma] - \mathcal{H}_{e,c}[\mu, \gamma]$ as:

$$\begin{aligned} & \left(\sum_{k=0}^e \binom{e}{k} \gamma^{(k)} \mu^{(e-k)} = \alpha^{(e)} \right) \\ & - \left(\sum_{k=0}^c \binom{e}{k} \gamma^{(k)} \mu^{(e-k)} = 0 \right). \end{aligned}$$

For $e \leq c$ we immediately obtain $\alpha^{(e)} = 0$ because the second summation goes only up to e since $\binom{e}{k}$ vanishes for $k > e$. For $e > c$ the relationship is rewritten as:

$$\left(\sum_{k=c+1}^e \binom{e}{k} \gamma^{(k)} \mu^{(e-k)} = \alpha^{(e)} \right)$$

Because $\gamma^{(c+1)} = 0$ this also leads to $\alpha^{(e)} = 0$. □

Proof of proposition 1, part II. We want to show the reverse implication, $\bar{\mathcal{R}}_{a,b}[\mu] \Leftrightarrow \mathcal{I}_{a,b}[\mu]$. We first prove the case $b > 0$. We use lemma 4, equation (14), which is $\mathcal{I}_{a,b}[\mu] \Rightarrow (\exists \gamma \in (C^\infty)^{b+1} \text{ s.t. } \hat{\mathcal{H}}_{a,b}[\mu, \gamma])$. We first show $\hat{\mathcal{H}}_{a,b}[\mu, \gamma]$ admits as solution (μ, γ) for which $\bar{\mathcal{R}}_{a,b}[\mu]$ holds and $\gamma = \gamma^{([0,b])}$. We then show this is the only solution to $\hat{\mathcal{H}}_{a,b}[\mu, \gamma]$. We have from lemma 4, equation (15), that $\hat{\mathcal{H}}_{a,b}[\mu, \gamma^{([0,b])}]$ is solvable. We have that $\hat{\mathcal{H}}_{a,b}[\mu, \gamma^{([0,b])}] \Leftrightarrow \hat{\mathcal{H}}_{a,b}[\mu, \gamma]$. By definition, $\hat{\mathcal{H}}_{a,b}[\mu, \gamma]$ implies $\mathcal{H}_{a+b,b}[\mu, \gamma] \wedge \mathcal{H}_{a+b+1,b}[\mu, \gamma]$ and we can thus apply lemma 5 with $c = b$ and $e = a + b$, leading to:

$$\hat{\mathcal{H}}_{a,b}[\mu, \gamma] \Rightarrow \gamma^{(b+1)} \mu^{(a)} = 0.$$

This implies that there exists $\Lambda \subset \mathbb{R}$ such that (i) $\mu^{(a)} = 0$ on Λ and (ii) $\gamma^{(b+1)}$ on $\mathbb{R} \setminus \Lambda$. In case (i), because μ is smooth from a premise of proposition 1, using equation (2) for μ in Λ implies $\mu \in \mathcal{P}_n$ with $n \leq a$, which is equivalent to $\bar{\mathcal{R}}_{a,0}[\mu]$. In case (ii), because γ is smooth from lemma 4, using equation (2) for γ in $\mathbb{R} \setminus \Lambda$ implies $\gamma \in \mathcal{P}_n$ with $n < b + 1$. We can use lemma 6 with $c = b$ and $e = a + 1$. Defining $\alpha = \gamma\mu$ we obtain that $\alpha^{(a+1)} = 0$, which implies through equation (2) that $\alpha \in \mathcal{P}_n$ with $n < a$. We can therefore write $\mu = \frac{\alpha}{\gamma}$ and thus $\bar{\mathcal{R}}_{a,b}[\mu]$. Combining cases (i) and (ii) we obtain that $\bar{\mathcal{R}}_{a,b}[\mu]$ holds since $\max(0, b) = b$. We finally prove the case $b = 0$ with $a \geq 0$. Using equation (6) we have that $\mathcal{I}_{a,0}[\mu] \Leftrightarrow \binom{a+1}{1} \mu^{(a+1)} = 0$. The factor $\binom{a+1}{1} = a + 1$ and because $a + 1 > 0$ we arrive at $\mu^{(a+1)} = 0$ and thus $\bar{\mathcal{R}}_{a,0}$ holds.

We now show that the above solution is the only one to $\hat{\mathcal{H}}_{a,b}[\mu, \gamma]$. We use an induction on b . The key observation to this proof is that for $b > 0$ the leading submatrix of $\beta_{a,b}[\mu]$ is $\beta_{a,b-1}[\mu]$. We start with $b = 0$. In that case $\beta_{a,0}[\mu] = \binom{a+1}{0} \mu^{(a+1)} = \mu^{(a+1)}$ where the last equality follows from $a + 1 > 0$. Because $\mathcal{I}_{a,0}[\mu] \Leftrightarrow (\mu^{(a+1)} = 0)$, as can easily be found from equation (6), we directly have that $\mu \in \bar{\mathcal{P}}_a$ which

implies $\bar{\mathcal{R}}_{a,0}[\mu]$. We now examine $b > 0$. In that case $\beta_{a,b}[\mu] \in C^\infty(\mathbb{R}, \mathbb{R}^{(b+1) \times (b+1)})$, and by definition of the invariant we have $\det(\beta_{a,b}[\mu]) = 0$ and so $\text{rank}(\beta_{a,b}[\mu]) \leq b$. We then face two cases. In the first case we have $\text{rank}(\beta_{a,b}[\mu]) < b$. This means that any $(b \times b)$ minor of $\beta_{a,b}[\mu]$ vanishes. In particular, deleting the last row and column of $\beta_{a,b}[\mu]$, we observe that the obtained submatrix is $\beta_{a,b-1}[\mu]$. This is easily understood from equation (13), where we observe that the entries of $\beta_{a,b}[\mu]$ do not depend on b (just its size depends on b). Therefore, from the induction hypothesis, we have $\bar{\mathcal{R}}_{a,b-1}[\mu]$ which implies $\bar{\mathcal{R}}_{a,b}[\mu]$. In the second case we have $\text{rank}(\beta_{a,b}[\mu]) = b$. The equations $\beta_{a,b}[\mu]\gamma = 0$ from $\hat{\mathcal{H}}_{a,b}[\mu, \gamma]$ are homogeneous. Any two solutions γ and γ' must thus be related by $\gamma = \tau\gamma'$ for some $\tau \in C^\infty(\mathbb{R}, \mathbb{R}^*)$. Therefore, we have that $\gamma = \tau\gamma^{([0,b])}$. Because $\hat{\mathcal{H}}[\mu, \gamma] = \hat{\mathcal{H}}[\mu, \tau\gamma^{([0,b])}] = \tau\hat{\mathcal{H}}[\mu, \gamma^{([0,b])}] = \tau\hat{\mathcal{H}}[\mu, \gamma]$, we conclude that all solutions to $\hat{\mathcal{H}}[\mu, \gamma]$ are equivalent. \square