Paraxial Geometric Optics in 3D through Point-based Geometric Algebra

Leo Dorst (l.dorst@uva.nl)

Computer Vision Group, Informatics Institute, University of Amsterdam, The Netherlands

Abstract. The versors of a homogeneous-point-based geometric algebra $\mathbb{R}_{d,0,1}$ (dubbed HGA) are related to the basic operations in geometric paraxial optics. Odd versors represent reflections in spherical mirrors (be they concave or convex) and even versors implement the lens equation. We extend the results to arbitrarily positioned optical elements by embedding $\mathbb{R}_{d,0,1}$ into CGA $\mathbb{R}_{d+1,1}$. The total transformation through a paraxial optical system now consists of successive teleportation (by CGA dot and outer product) to the next optical center, and then applying its local HGA versors.

The result is a straightforward sequence of operations which implements a total system of arbitrarily placed paraxial lenses and mirrors in 3D (or any dimension), parameterized by their CGA tangent vectors (from each optical center to the corresponding focal point) for each optical component. This can be used to compile the homogeneous transformation matrices of a total paraxial system in terms of those geometric parameters.

A similar talk was given at CGI/ENGAGE 2023, and published as [1]. It is submitted to AGACSE 2024 to instruct the GA community on the meaning of the 'dual PGA'.

1 HGA: The Algebra of Homogeneous Point Coordinates

Ray transfer matrices have traditionally been used to compute with planar paraxial optical systems, but in a height/slope parametrization of rays that includes a needless linearizing approximation. Recently, [2] showed that by using homogeneous coordinates in 2D, the ray matrices can be exact, and matrices for point imaging also be included. The homogeneous matrices of rigid body transformations can then be employed to process optical systems with different optical axes, still on the 2D optical table. In the present paper, we demonstrate how the geometric algebra HGA of homogeneous coordinates affords a natural parametrization to unify the imaging of geometric primitives by a generally placed system of paraxial optical elements in 3D space (with 2D still included, of course). It can then be used to generate the corresponding 4×4 homogeneous matrices, if desired.

The homogeneous coordinates of a point at location $\mathbf{x} = [x_1, \cdots, x_d]^\top$ in a Euclidean space $\mathbb{R}^d$ are obtained by adding one extra representational dimension; they are $[1, x_1, \cdots, x_d]^\top$. In geometric algebra, we introduce a basis vector e_0 for the extra dimension, and we need to decide the metric relationships for all vectors. The metric for the Euclidean part remains Euclidean (so $\mathbf{e}_i \cdot \mathbf{e}_j = \delta_{ij}$ for an orthonormal basis). Relative to the Euclidean basis vectors $\mathbf{e}_i$, we set $e_0 \cdot \mathbf{e}_i = 0$; the extra representational dimension is orthogonal to the Euclidean spatial aspects. But we will explicitly choose e_0 to be a *null vector* (i.e., $e_0^2 = 0$), so that our algebra is of signature $\mathbb{R}_{d,0,1}$; let us call it HGA, for *homogeneous geometric algebra.*[1]

[1] This signature differs from the 'homogeneous model' in [3] (where $e_0^2 = \pm 1$), but is used in [4] for the point-based algebra of Euclidean space: not to be confused with plane-based Euclidean PGA $\mathbb{R}^*_{d,0,1}$ [5,6].

1.1 Applying the Vectors of HGA $\mathbb{R}_{d,0,1}$

HGA $\mathbb{R}_{d,0,1}$ has three kinds of vectors, differing in their algebraic properties and their geometric semantics: purely Euclidean vectors $\mathbf{m}$ denoting directions, the (only!) null vector e_0 denoting the point at the origin, and a general point vector $R = e_0 + \mathbf{r}$. Only e_0 can not be used as a reflector.

1.2 Reflection in a Direction: Planar Mirror

A purely Euclidean vector indicates a 1-dimensional direction. Such a vector versor $\mathbf{m}$ acts on a point $X = e_0 + \mathbf{x}$ by sandwiching, to produce:

$$-\mathbf{m}\,X\mathbf{m}^{-1} = -\mathbf{m}\,(e_0 + \mathbf{x})\,\mathbf{m}^{-1} = e_0 - \mathbf{m}\mathbf{x}\mathbf{m}^{-1},$$

and we recognize in $-\mathbf{m}\mathbf{x}\mathbf{m}^{-1}$ the reflection of the vector $\mathbf{x}$ in a plane with normal $\mathbf{m}$ passing through the origin. Thus X reflects in the origin plane with normal $\mathbf{m}$; the point is seen at the other side of the mirror, at the same perpendicular distance as X (see Figure 1a).

1.3 Reflection in a Point: Spherical Mirror

A general unit weight vector $R = e_0 + \mathbf{r}$ (geometrically the point at location $\mathbf{r}$) can also be used in sandwiching as a versor to transform points. With $R^{-1} = R/\mathbf{r}^2$ we obtain:

$$-R\,(e_0 + \mathbf{x})/R = \;\cdots\; = (1 - 2\mathbf{r}^{-1}\cdot\mathbf{x})\,\Big(e_0 + \frac{-\mathbf{r}\mathbf{x}\mathbf{r}^{-1}}{1 - 2\mathbf{r}^{-1}\cdot\mathbf{x}}\Big). \tag{1}$$

Note in Equation 1 how the versor action produces an additional factor for e_0, proportional to the $\mathbf{r}$-component of the input term $\mathbf{x}$. In the final expression, we factored out the 'weight' of the point to expose its Euclidean location. The location is a vector proportional to the reflection of $\mathbf{x}$ in the plane with normal $\mathbf{r}$ by an $\mathbf{x}$-dependent factor $1/(1 - 2\mathbf{r}^{-1}\cdot\mathbf{x})$. If $\mathbf{r}^{-1}\cdot\mathbf{x}$ is sufficiently large, that factor is negative, and the image is a negatively weighted point at the same side of the $\mathbf{r}$-plane as the input point $\mathbf{x}$.

When we take $\mathbf{r}$ on the negative side of the origin on the optical axis (in the Cartesian sign convention), this is the GA form of the formula for *reflection in a concave spherical mirror* with spherical center $R = e_0 + \mathbf{r}$, in the paraxial approximation of geometric optics, see Figure 1b.

1.4 The Lensing Versor $L_{\mathbf{f}}$ as Double Reflection

It is now natural to consider a lens as the combination of a spherical mirror and a reflection, so to use a versor $L = \mathbf{r}R$ to represent it. After some manipulation setting $\mathbf{f} = -\mathbf{r}/2$ this gives:

$$L_{\mathbf{f}} \equiv 1 + \tfrac{1}{2}e_0\mathbf{f}^{-1},$$

which used in sandwiching results in the homogeneous point

$$L_{\mathbf{f}}\,(e_0 + \mathbf{x})\,L_{\mathbf{f}}^{-1} = \cdots = (1 + \mathbf{f}^{-1}\cdot\mathbf{x})\,\Big(e_0 + \frac{\mathbf{x}}{1 + \mathbf{f}^{-1}\cdot\mathbf{x}}\Big). \tag{2}$$

The transformed location indeed corresponds to the classical result of lensing, see Figure 1c.

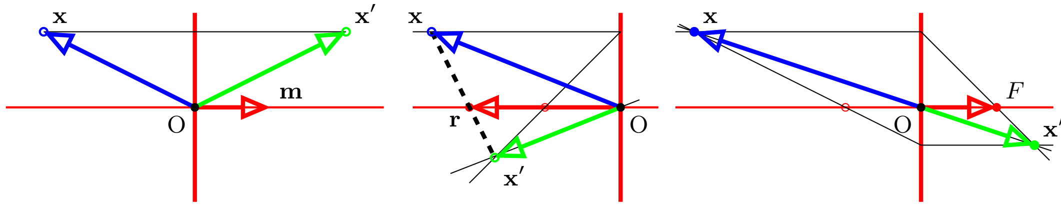

(a) (b) (c)

Fig. 1. (a) The reflection in a planar mirror represented by a vector $\mathbf{m}$; here $\mathbf{x}$ points at the input point, $\mathbf{x}'$ at the output point. (b) Paraxial reflection in a spherical mirror with optical center at O and radial center at $\mathbf{r}$ is represented as versor sandwiching by the homogeneous point at $\mathbf{r}$. It transforms directly along the dotted line, or by the usual ray construction using a focal point halfway. (c) A convex lens wth focal point F in the paraxial approximation of geometric optics, viewed in a plane containing the optical axis.

1.5 Imaging Arbitrary Flats; Homogeneous Matrix Representation

Once we have specified how points are imaged, we can of course also know how lines and planes are imaged. This is true in any formalism, but in HGA it takes a particularly simple form, due to the versor nature of the paraxial mappings. Applying the lensing versor, a general flat with directional part $\mathbf{A}$ and passing through $\mathbf{p}$ lenses to:

$$(e_0 + \mathbf{p}) \wedge \mathbf{A} \overset{L_\mathbf{f}}{\mapsto} e_0 \wedge \big(\mathbf{A} + \mathbf{f}^{-1} \cdot (\mathbf{p} \wedge \mathbf{A})\big) + \mathbf{p} \wedge \mathbf{A}. \tag{3}$$

If you wish, you could compose a *matrix* for the linear map $X \mapsto L_\mathbf{f} X / L_\mathbf{f}$, rather than characterize it by a versor. However, this matrix would depend on the type of element X (just like the 2D homogeneous matrices from [2]). Using the classical homogeneous characterizations of points, lines and planes, one then obtains the matrices

$$\begin{bmatrix}\mathbf{p}\\1\end{bmatrix} \overset{L_\mathbf{f}}{\mapsto} \begin{bmatrix}[1] & \mathbf{0}\\ [\mathbf{f}^{-1}]^\top & 1\end{bmatrix}\begin{bmatrix}\mathbf{p}\\1\end{bmatrix} \text{ and } \begin{bmatrix}\mathbf{u}\\ \mathbf{p}\times\mathbf{u}\end{bmatrix} \overset{L_\mathbf{f}}{\mapsto} \begin{bmatrix}[1] & \mathbf{0}\\ [\mathbf{f}^{-1}]^\times & 1\end{bmatrix}\begin{bmatrix}\mathbf{u}\\ \mathbf{p}\times\mathbf{u}\end{bmatrix} \text{ and } \begin{bmatrix}\mathbf{m}\\ -\delta\end{bmatrix} \overset{L_\mathbf{f}}{\mapsto} \begin{bmatrix}[1] & -[\mathbf{f}^{-1}]\\ \mathbf{0}^\top & 1\end{bmatrix}\begin{bmatrix}\mathbf{m}\\ -\delta\end{bmatrix}.$$

2 Paraxial Geometric Optics Anywhere in Space

The null vector e_0 encoding the optical center makes the lensing versor work properly. If we want to move the lens to another location than the origin, we therefore need a null vector at that new location. That is not possible in HGA, where e_0 is the only null point. Rigid body motions are simply not among the versors of HGA, all of which we just exposed.

2.1 Placing a HGA at Any Spatial Location

We move to the larger algebra CGA $\mathbb{R}_{d+1,1}$ (conformal geometric algebra, see e.g. [3]), which has *all* points as null vectors, and construct a copy of HGA at every point. Then we can perform the earlier lens versor construction at any point we like. Concatenation of lenses will be done by hopping from one optical center to the next, for each choosing the corresponding algebraic embedding.

A point at location $\mathbf{p}$ in CGA is represented by the vector $p = o + \mathbf{p} + \frac{1}{2}\mathbf{p}^2\infty$; this is a null vector. The origin point of HGA denoted by its null vector e_0 seems to correspond naturally to the vector o of CGA; however, the whole purpose of embedding into CGA is our desire to have a HGA at *any* point p of CGA. Since any CGA point is a null vector (a sphere of zero radius), such recasting will not affect the local algebra and geometry of lensing as designed with HGA.

A geometrical point (such as might occur in the intersection of a line and a plane) is a 'flat point' in CGA, of the form $p \wedge \infty$, since those intersecting flat elements always also have the point at infinity ∞ in common. A unit flat point squares to 1: $(p \cdot \infty)^2 = (\infty \cdot p)^2 = 1$. The flat point can be rewritten as $p \wedge \infty = (o + \mathbf{p}) \wedge \infty$, and we recognize in the first factor a natural identification with the point representation $e_0 + \mathbf{p}$ of HGA. With the identification of the origin elements of the two algebras $e_0 = o$, we can write the HGA point $e_0 + \mathbf{p}$ as an element $o + \mathbf{p}$ of CGA,[2] parametrized by CGA point p:

$$e_0 + \mathbf{p} \leftrightarrow o + \mathbf{p} = o \cdot (-\infty \wedge p). \tag{4}$$

This final form of the point representation contains explicitly the point o which we took as the origin of our local HGA. Without changing the algebra, we can now choose any other CGA point as corresponding to the null vector e_0 of HGA; then a HGA point at location $\mathbf{x}$ but viewed from another 'origin' c is represented by the CGA element:

$$x|_c \equiv c \cdot (-\infty \wedge x). \tag{5}$$

where c and x are CGA points. You may pronounce '$|_c$' as 'from c'. This operation considers x in a copy of HGA at the location c. It is structure-preserving: the formation of new elements by the HGA outer product is equivariantly preserved by this construction (it is a *linear outermorphism*):

$$x|_c \wedge y|_c = \big(c \cdot (-\infty \wedge x)\big) \wedge \big(c \cdot (-\infty \wedge y)\big) = \;\cdots\; = c \cdot (-\infty \wedge x \wedge y) = (x \wedge y)|_c.$$

Therefore an arbitrary HGA element X (point, line, plane, direction element) can be embedded into CGA as a 'from c'-element through:

$$X \;\mapsto\; X|_c \equiv c \cdot (-\infty \wedge X), \tag{6}$$

(where we substitute HGA's e_0 always by CGA's o before putting X in the 'from c' formula, to make the formula computable within CGA). Note that $c|_c = c$.

The original HGA element X can be retrieved from this as:

$$X|_c \;\mapsto\; X = o \cdot \big(-\infty \wedge (X|_c)\big), \tag{7}$$

and then substituting $e_0 \mapsto o$ (to formally get out of CGA back to standard HGA).

A highly useful property is that the 'from c' mapping can be applied multiple times, but that the result only depends on the *last* application (we could call this '*neopotent*'):

$$(X|_{c_1})|_{c_2} = X|_{c_2}. \tag{8}$$

We can therefore always rerepresent a rerepresented element, and in a concatenation of operations there is no need to revert from $X|_c$ to the original X before we can perform the next step. Jumping to a new viewpoint c is not a relative translation, but an absolute teleportation.

[2] Considered as an element of CGA, $o \cdot (-\infty \wedge p)$ is geometrically an oriented dual sphere with center p, and passing through o. While that geometric interpretation can be maintained through the subsequent algebra, it is rather distracting, so we will not emphasize it.

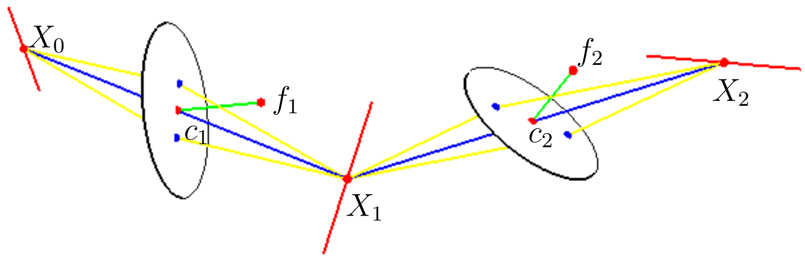

Fig. 2. Propagation of a point on a line through two lenses parametrized by centers and foci in 3D using Equation 13 implemented in GAviewer [8].

2.2 The HGA Lens Versor in CGA

For a lens with optical center c, we should move its lens versor (which was $\exp(\frac{1}{2}e_0 \wedge \mathbf{f}^{-1})$ at the origin) also to that location. With the substitution $e_0 \mapsto o$, the lens versor becomes a versor $\exp(\frac{1}{2}o \wedge \mathbf{f}^{-1})$ in CGA (it is a conformal transformation called a *transversion*, see [3]). It involves the CGA tangent vector $o \wedge \mathbf{f}^{-1}$, and it and its versor can be moved to c by the CGA translation versor $\exp(-\infty \wedge \mathbf{c}/2)$. The result can be rewritten into the form (see [1]):

$$\ell_{c,f} = c/(c \wedge \infty \wedge f) = c/(c \wedge \infty \wedge \mathbf{f}), \tag{9}$$

with f the focal point, and $\mathbf{f}$ the relative vector from c to f. Note that in this format, either point f or relative vector $\mathbf{f}$ could be used as input parameter for an identically defined function computing the tangent vector; no conversion is required since $c \wedge \infty \wedge f = c \wedge \infty \wedge \mathbf{f}$. With that tangent vector $\ell_{c,f}$, the lens versor is: $L_{c,f} = \exp(\frac{1}{2}\ell_{c,f}) = 1 + \frac{1}{2}\ell_{c,f}$.

The full lens mapping cannot simply be this CGA versor $L_{c,f}$ applied to a CGA point – that would be a conformal transformation, and lensing is not (it transforms circles to ellipses, not to other circles). In our local copy of HGA at the point c, the lens versor should act on an element X by first converting that to the 'relative to c' form $X|_c$ and then applying the c-based versor:

$$X \quad \mapsto \quad X'|_c = L_{c,f}\,(X|_c)\,L_{c,f}^{-1} \equiv \underline{L_{c,f}}[X|_c]. \tag{10}$$

(The underline notation is a common compact way to denote the versor sandwiching/conjugation operation.) The result $X'|_c$ is again a 'from c' type element of CGA. The corresponding flat element is $-\infty \wedge (X'|_c)$, but it is more convenient to keep it in the 'from c' form if one wants to apply another lens to it next. We can write Equation 10 more explicitly, and rewrite slightly, using $\underline{L_{c,f}}[c] = c$:

$$c \cdot (-\infty \wedge X') = \underline{L_{c,f}}[c \cdot (-\infty \wedge X)] = c \cdot \underline{L_{c,f}}[-\infty \wedge X]. \tag{11}$$

The latter shows clearly that the lens versor action is concentrated on the flat elements of CGA.

The spherical mirroring versor $R = o - 2\mathbf{f}$ can also be brought into a form in which it may be parametrized by either the relative Euclidean vector $\mathbf{f}$ or by the CGA point f, namely: $R = c - 2\,(c \wedge \infty \wedge f)/(c \wedge \infty)$. Remember from general GA that this odd versor should be applied to an element $X|_c$ as $\underline{R}[X|_c] = R\,\widehat{X}|_c\,R^{-1}$, the grade involution $\widehat{\ }$ giving a minus sign for odd-grade X.

2.3 Concatenation of Lenses

With the above, one can compute the paraxial image of an element X in HGA through a succession of n lenses and/or mirrors in d-dimensional space, by the CGA embedding.

1. Let the lenses have optical centers at CGA points c_i, and focal points at CGA points f_i (or have relative focal vectors $\mathbf{f}_i$ from their c_i). Form the *lens versors*
$$L_i \equiv 1 + \tfrac{1}{2}\, c_i/(c_i \wedge \infty \wedge f_i) \tag{12}$$
or the *spherical mirror versor* $R_i = c_i - 2(c_i \wedge \infty) \cdot (c_i \wedge \infty \wedge f_i)$, or the *planar mirror versor* $M_i = (c_i \wedge \infty) \cdot (c_i \wedge \infty \wedge \mathbf{f}_i)$ (for which $\mathbf{f}_i$ is the mirror normal vector pointing from c_i to f_i).
2. Embed the HGA element X into CGA by replacing its e_0 by o. Then perform the iteration:
$$X_0 = X, \quad X_i = c_i \cdot \underline{L_i}[-\infty \wedge X_{i-1}] \quad \text{for} \quad i = 1, \cdots, n. \tag{13}$$
(or similarly for $\underline{R_i}[]$ and $\underline{M_i}[]$, with the grade involution in both sandwichings).
3. After processing all n optical elements, the result is $X' = o \cdot (-\infty \wedge X_n)$ relative to an origin o; it can be converted back to HGA by replacing o by e_0, if desired.

Figure 2 shows a 3D example of a point and line imaged through two lenses by this method.

2.4 Generating Optical System Matrices

Since any flat geometric primitive can be propagated through the system, it is now easy to find the total homogeneous matrix for a composition of optical elements, for any flat geometric element. For instance, if you need the matrix for the imaging of an arbitrary 3D line, use the Plücker coordinate basis $\{\mathbf{e}_{01}, \mathbf{e}_{02}, \mathbf{e}_{03}, \mathbf{e}_{23}, \mathbf{e}_{31}, \mathbf{e}_{12}\}$ to represent both it and the result. Simply process the i-th basis element by Equation 13 and denote the resulting components as the i-th column of the transformation matrix. This extends the 2D techniques of [2] to 3D, and conveniently parametrizes the system by the absolute position and focal points of the optical elements.

References

1. L. Dorst, "Paraxial geometric optics in 3d through point-based geometric algebra," in *CGI 2023, LNCS 14498*, B. e. a. Sheng, Ed. Springer Verlag, 2024, pp. 340–354.
2. T. Corcovilos, "Beyond the ABCDs: A better matrix method for geometric optics by using homogeneous coordinates," *American Journal of Physics*, vol. 91, pp. 449–457, 2023.
3. L. Dorst, D. Fontijne, and S. Mann, *Geometric Algebra for Computer Science: An Object-oriented Approach to Geometry*. Morgan Kaufman, 2009.
4. E. Lengyel, 2022. [Online]. Available: https://projectivegeometricalgebra.org/
5. C. Gunn, "On the homogeneous model of Euclidean geometry," in *Guide to Geometric in Practice*, L. Dorst and J. Lasenby, Eds. Springer-Verlag, 2011, pp. 297–327. [Online]. Available: https://www.springer.com/gp/book/9780857298102
6. ——, "Geometric algebras for Euclidean geometry," *Advances in Applied Clifford Algebras*, vol. 27, pp. 185–208, 2017.
7. Q. Sugon and D. McNamara, "Paraxial meridional ray tracing equations from the unified reflection-refraction law via geometric algebra," 2008. [Online]. Available: https://arxiv.org/abs/0810.5224
8. L. Dorst and D. Fontijne, "3D Euclidean geometry through conformal geometric algebra (a GAViewer tutorial)," `http://www.science.uva.nl/ga/`.