Reconstruction and Visualization of Planetary Nebulae

Marcus Magnor, *Member*, *IEEE*, Gordon Kindlmann,
Charles Hansen, *Member*, *IEEE*, and Neb Duric

Abstract—From our terrestrially confined viewpoint, the actual three-dimensional shape of distant astronomical objects is, in general, very challenging to determine. For one class of astronomical objects, however, spatial structure can be recovered from conventional 2D images alone. So-called planetary nebulae (PNe) exhibit pronounced symmetry characteristics that come about due to fundamental physical processes. Making use of this symmetry constraint, we present a technique to automatically recover the axisymmetric structure of many planetary nebulae from photographs. With GPU-based volume rendering driving a nonlinear optimization, we estimate the nebula's local emission density as a function of its radial and axial coordinates and we recover the orientation of the nebula relative to Earth. The optimization refines the nebula model and its orientation by minimizing the differences between the rendered image and the original astronomical image. The resulting model allows creating realistic 3D visualizations of these nebulae, for example, for planetarium shows and other educational purposes. In addition, the recovered spatial distribution of the emissive gas can help astrophysicists gain deeper insight into the formation processes of planetary nebulae.

Index Terms—Astronomy, constrained optimization, volumetric image representation, scene analysis, volume visualization.

1 Introduction

Due to their diverse and colorful appearance, planetary nebulae (PNe)[1] are among the most esthetically attractive objects in the night sky. Pictures of PNe are regularly published in public science magazines and daily newspapers as their visual appeal make PNe ideal objects to stir broad public interest in astrophysics research. PNe even concern us directly: Our own sun will end its life as a planetary nebula, albeit not until some four billion years from now.

Planetary nebulae become still more vivid if they can be realistically visualized in 3D. For planetarium shows or TV documentaries, being able to authentically visualize PNe enables creating impressive animations, such as virtual fly-arounds, to captivate the audience as well as to illustrate the underlying physics. But, visualizing the three-dimensional shape of PNe can also be of scientific value to astrophysicists, who may use PN visualization to investigate the gas distribution in space and to evaluate simulation results derived from theoretical nebula models. Our paper describes a computer graphics-based approach to how to simultaneously recover and realistically visualize the 3D shape of planetary nebulae from conventional photographs.

While it was soon realized that what we observe of PNe are 2D projections of 3D volumes of glowing gas, the actual spatial structure of PNe has long been a subject of investigation in astrophysics. It was not until the second half of the 20th century that, based on earlier morphological classifications [1], qualitative models were proposed to explain the general appearance of PNe [2], [3], [4]. While these simple models succeeded in qualitatively describing the great variety of observed PN shapes, they were not accurate enough to reflect the actual three-dimensional structure of individual planetary nebulae.

To recover the actual structure of planetary nebulae from conventional 2D images, we propose a volume modeling technique coined *constrained inverse volume rendering* (CIVR). Based on fast volume rendering using PC graphics hardware, we synthesize realistic views from a generic PN model. This model incorporates the symmetry constraints arising from the physical formation processes of bipolar PNe by representing the spatial shape as an axisymmetric emission density distribution map. An optimization algorithm varies the model's emission distribution entries and iteratively compares the rendered model to the actual image until the rendered model best matches the observed image. By exploiting the axisymmetric nature of most PNe to constrain the reconstruction process, the inverse problem becomes well-posed and, at the same time, we obtain only physically plausible results. Reconstruction and visualization of PNe is useful for the production of scientifically consistent animations for planetarium shows and TV documentaries, as well as for the development of interactive tools for educational purposes. In astrophysics research, the recovered spatial distribution of the emissive gas can aid in gaining a deeper understanding of PN

1. We will use "PN" and "PNe" to refer to planetary nebulae in the singular and plural, respectively.

- *M. Magnor is with the MPI Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany. E-mail: magnor@mpi-sb.mpg.de.*
- *G. Kindlmann is with the Surgical Planning Laboratory, Brigham and Women's Hospital, Boston, MA 02115. E-mail: gk@cs.utah.edu.*
- *C. Hansen is with the School of Computing, University of Utah, Salt Lake City, UT 84112. E-mail: hansen@cs.utah.edu.*
- *N. Duric is with the Department of Physics and Astronomy, University of New Mexico, Albuquerque, NM 87106. E-mail: duric@unm.edu.*

Manuscript received 1 Nov. 2004; revised 21 Jan. 2005; accepted 27 Jan. 2005; published online 11 July 2005.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCG-Magnor Hansen-1004.

formation processes. Reconstructed PN shapes can, for example, be compared to computer simulation results to validate theoretical models of the stellar wind. A database of PN shapes may also help in establishing a general timeline of PN evolution.

In this paper, we extend our earlier work on the topic [5]. We review previous related work in Section 2. After an introductory overview of the fundamental physics of PNe in Section 3, we describe our constrained inverse volume rendering approach in Section 4. Section 5 illustrates the practical steps in reconstructing planetary nebulae. Results are presented in Section 6. We discuss our method and validate its performance in Section 7 before we conclude in Section 8.

2 Related Work

The presented volumetric modeling technique relies on recent advances in visualization and graphics research. The term *inverse volume rendering* was first employed by Curless and Levoy [6] to refer to updating volume data from range images. Mueller et al. [7] use the same term to refer to a reprojection technique in CT reconstruction. Our use of the term is derived from Marschner's PhD thesis [8] in which he introduces the term *inverse rendering* to denote the idea of reconstructing scene properties from image data via computer graphics techniques.

The volume modeling approach we present is based on the ability to perform *volume rendering* [9] very efficiently on modern graphics hardware. Cabral et al. [10] demonstrate how to perform efficient volume rendering using texture mapping hardware, a fundamental technique that is still applicable to today's programmable graphics architectures [11]. For purely self-emitting volumes, Cabral et al.'s slicing technique directly yields correct rendering results [12].

While spectacular 3D fly-throughs of astronomical objects can be experienced in many planetariums, these animations are almost always purely artistic pieces of work. One notable exception is the effort by Nadeau et al. [13], who employed massive computational power to create scientifically justified views of the Orion nebula for the Hayden Planetarium in New York [14]. For their visualizations, they rely on a 3D model of the Orion nebula that was determined by astronomers from decades of various observational data [15]. Our method, in contrast, simultaneously reconstructs the 3D volume model from photographs and visualizes it realistically, albeit for a different class of astronomical objects: planetary nebulae.

In astrophysics, a large body of literature concerns the morphology and shape of PNe. Curtis [1] classified the planetary nebulae known at his time into different classes based on their visual appearance. Khromov and Kohoutek [2] qualitatively described the diversity in PN appearance as an effect of nebula orientation from one generic 3D shape. Balick [3], Kwok et al. [16], and Calvet and Peimbert [17] introduced the interacting-winds theory to explain the physical processes leading to the 3D structure of PNe. Making use of involved astronomical observation techniques, 3D gas velocities have been directly measured and the three-dimensional gas distribution inferred for a handful of PNe by Sabbadin et al. [18], [19] and Saurer [20]. To investigate the physical processes leading to the observed PN structures, gas-kinetic and photo-ionization simulations have been conducted by, among others, Mellema and Frank [21], Armsdorfer et al. [22], and Ercolano et al. [23]. The simulation results qualitatively confirm observational findings. Limited computational capacity, however, currently prevents the simulation parameters from being driven by observational data and visual correspondence between simulation results and individual planetary nebulae has yet to be established.

3 Planetary Nebulae

In the following, we give a brief explanatory overview of the main physical processes that lead to the visual appearance of PNe. For a more thorough introduction to the physics of gaseous nebulae and PNe, we recommend the classic books on the topic by Osterbrock [24] and Kwok [25]. A concise overview of our current knowledge about the formation processes leading to planetary nebulae has recently been published by Balick and Frank [26].

Despite their name, PNe are of extra-solar origin: A planetary nebula constitutes the last episode in the life of any star weighing less than about eight sun masses. At this stage, the star has burned up most of its nuclear fuel. On astronomical scales, PNe are extremely short-lived, existing only for a few tens of thousands of years before they fade away, which is why only about 1,500 have been discovered in our Galaxy so far. Owing to their great distances from Earth, they are dim and small, spanning at most a few arc minutes in the sky. For the work presented here, we rely on publicly available images of PNe that have been recorded with the Hubble Space Telescope [27] and the Nordic Optical Telescope at Teneriffe, Canary Islands [28].

3.1 Photo-Ionization

When a star of initially less than eight times the mass of the Sun nears its end of energy production, a strong wind blows off its surface, carrying substantial amounts of stellar matter into space and exposing deeper, hotter regions of the star. This wind consists mainly of hydrogen, but also of other elements that were previously synthesized within the star's core. At the same time, the central star contracts, which increases the exposed surface's temperature further. Following Wien's law for blackbody radiation [24], with rising surface temperature, the star emits more high-energy ultra-violet (UV) photons. Once the star emits enough UV photons carrying sufficient energy to ionize hydrogen (13.6 eV, corresponding to a wavelength of 91nm), the ejected hydrogen gas becomes completely photo-ionized, i.e., essentially all hydrogen atoms are being separated into free electrons and hydrogen ions (protons). Within this plasma, both constituents frequently recombine to form neutral hydrogen atoms again. These recombined atoms are typically in an excited state, however, from which they return to the ground state via a cascade of energy transitions before another UV photon ionizes the atom again. Each transition during this cascade is accompanied by the emission of a photon of a distinct wavelength. These emission-line photons can escape from the nebula essentially unhindered and it is

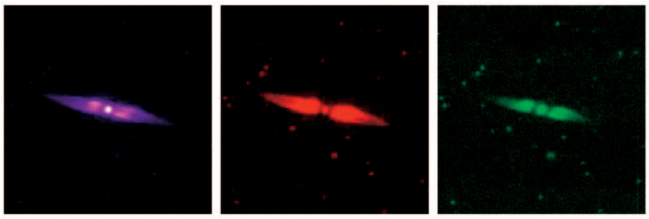

Fig. 1. Input images: The planetary nebula He2-437 in the light of hydrogen, nitrogen, and oxygen (from left to right). While the original image data is scalar-valued (gray-scale images), correct color appearance is determined by taking into account the complete emission line spectrum of each chemical element (Section 6.1).

these photons that we observe and that are responsible for the nebula's characteristic color.

Besides hydrogen, the other chemical elements in the nebula are also in ionized states and emit photons at characteristic wavelengths. Thus, images recorded at different wavelengths enable astronomers to identify different elements in different regions of a nebula, as shown in Fig. 1. For most (mature) PNe, the central star's UV photon flux is more than sufficient to photo-ionize all gas that had previously blown off the star. Assuming that the amount of light-absorbing dust can be neglected, we observe the entire mass of ejected gas as one glowing volume in space [25].

It must be noted that local gas volume emissivity, i.e., the number of photons emitted from a unit volume per unit time, is actually proportional to the square of local gas density. This is because the statistical probability for electron-ion recombinations per unit time is proportional to the number of hydrogen ions per unit volume times the number of electrons per unit volume, both numbers being approximately the same. We deliberately do not make this distinction between local photon density and gas density in order to avoid further complicating the following text passages.

3.2 Axisymmetry

Planetary nebulae glow because gas previously ejected from the central star is being ionized by the star's intense UV radiation. Photo-ionization alone, however, cannot explain the axial symmetry of most PNe since the central star certainly radiates isotropically into all directions. Instead, it is the nonuniform spatial distribution of the gas that is responsible for the axisymmetric appearance of many PNe.

A classification scheme for PNe based on visual morphology was proposed by Curtis early in the 20th century [1]. In the 1960s, it became clear that the various different shapes of PNe may be qualitatively explained by one common, general shape viewed from Earth at different orientations [2]. While this phenomenological model was motivated by observational data, it could not explain the physical processes leading to the shape. Kwok et al. [16] and Balick [3] finally proposed a mechanism that explains the observed gas distribution within PNe, illustrated in Fig. 2; their *interacting stellar wind theory* starts with the premise that, prior to its final stage, the star has already been blowing off gas, but at a much slower speed. This continuously ejected gas accumulates in the star's equatorial plane. Toward the end of the star's life, however, the velocity of the discharged

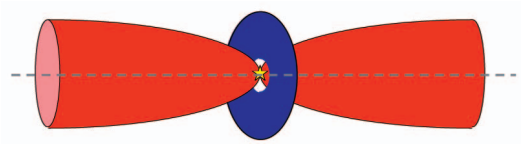

Fig. 2. Interacting winds theory: During its lifetime, a torus of dense, slowly outward moving gas gathers in the equatorial plane of the central star (blue). At the end of its life, the star suddenly blows off gas at a much higher rate. When this fast wind hits the inner boundary of the old gas masses, the fast-moving gas is redirected and escapes in opposite, polar directions.

gas increases by several orders of magnitude. Now, this much faster and hotter wind collides with the old gas masses within the star's equatorial plane. The old, slow gas masses act like a nozzle and the fast wind is redirected away from the equator and escapes toward the poles. Two lobes of gas form, flowing away from the central star in opposite directions. From the overall symmetry of the system and the conservation of momentum, the outflowing gas jets must be axisymmetric, creating the axisymmetric shape of many PNe.

Besides these interacting winds [29], recent research suggests that, in the case of tight binary star systems, the companion star plays an important role in the formation of the planetary nebula [26]. Finally, the central star's magnetic field can additionally focus the outflow of the photo-ionized, charged gas plasma, as is probably the case for our test nebulae He2-437, depicted in Fig. 9.

Analytical calculations [30] as well as extensive gas-dynamic [21] and photo-ionization simulations [31], [23] demonstrated the plausibility of the interacting stellar wind theory, while observational evidence for its correctness was first put forward by Balick [3]. On its basis, a low-dimensional (4 degrees-of-freedom) shape model for PNe, the *prolate ellipsoidal shell* (PES) model, was proposed by Masson [4]. The PES model was subsequently used by Aaquist and Kwok [32] and Zhang and Kwok [33] to model the general appearance of a number of PNe by adjusting the four model parameters by hand. This low-dimensional model, however, can only give a very rough, qualitative account of the spatial gas density distribution within PNe. For realistic visualization purposes, as well as to investigate the gas outflow characteristics of actual PNe, a much more detailed model description must be recovered.

4 Constrained Inverse Volume Rendering

The physics of PNe enables us to develop a method to reconstruct the axisymmetric emissive gas distribution from two-dimensional images as observed from Earth. Our *constrained inverse volume rendering* (CIVR) approach relies on three key physical and observational properties:

- Most PNe exhibit a symmetry axis. This axisymmetry reduces the 3D volumetric reconstruction problem to a 2D emission density function in cylindrical coordinates.
- Scattering and absorption is assumed to be negligible at visible wavelengths at which our astronomical input images are acquired.

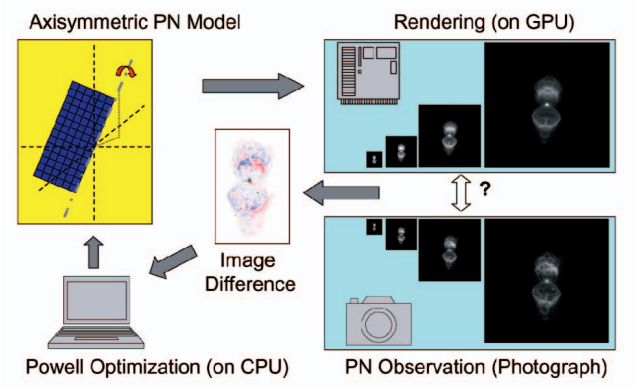

Fig. 3. Constrained Inverse Volume Rendering (CIVR) workflow.

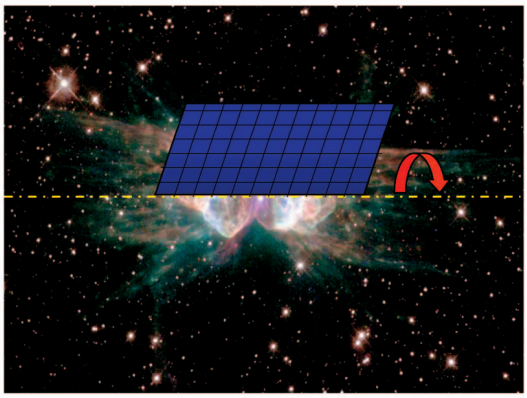

Fig. 4. Axial symmetry: The structure of most planetary nebulae can be described by a density map in cylindrical coordinates. The volumetric density distribution in 3D space is generated by rotating the density map along the nebula's axis of symmetry.

- Planetary nebulae span only a few arc minutes in the sky (less than a tenth of a degree), so their projected appearance is essentially orthographic.

Our CIVR algorithm iteratively determines PN model parameter values until the rendered model image best matches the nebula's photograph. Fig. 3 illustrates the algorithm:

- The PN model is realistically rendered on the GPU.
- The rendered image is compared to the actual photograph.
- The error function is evaluated on the CPU.
- The error function value drives a nonlinear optimization algorithm that
- varies PN model parameter values prior to the next iteration step.

The PN model consists of a two-dimensional density map whose entries represent axisymmetric emission density values. As depicted in Fig. 4, the map is aligned with the symmetry axis of the nebula such that the central star is in the middle. The emission map entries and the two orientation angles of the symmetry axis represent the model parameters that need to be optimized. Making use of a modern PC graphics board, the projection image of the PN model is rendered. The rendered image is compared to the actual photograph of the PN by calculating per-pixel differences. A scalar error value is derived that drives a nonlinear optimization algorithm running on the CPU. The optimization routine, in turn, varies the parameter values of the PN model such that the error function value decreases. To accelerate convergence, a hierarchical optimization approach is pursued, based on an image resolution pyramid. The CIVR algorithm terminates either when the error function value levels out, when it reaches a preset minimum threshold, or when a preset number of maximum iterations has been computed.

4.1 Emissive Volume Rendering

Because CIVR can take up to 10^6 iterations to converge and rendering the PN model is by far the most time-consuming operation, we must be able to render the PN model as fast as possible without sacrificing physical accuracy. Since PNe are made of glowing gas, synthesizing faithful projection images from our PN model is equivalent to rendering arbitrary views of a purely emissive volume. Using ray tracing techniques, such emissive volume rendering can be straightforwardly implemented on the CPU. However, the permanent transformations from viewpoint coordinates to the oriented model's frame of reference and on to the cylindrical coordinates of the emissive density map slows down software rendering performance considerably.

Alternatively, emissive volume rendering can be implemented on modern PC graphics hardware. Multiple viewport-aligned slices through the nebula volume are rendered at increasing distances from the viewpoint, while all slices' contributions are accumulated [10]. Each quadrilateral sent through the graphics pipeline spans the entire viewport and is oriented perpendicular to the viewing direction. The transformation from view to model coordinates is done at no cost by loading the appropriate model view matrix to the stack. Programmable fragment processing then enables us to convert on-the-fly rasterized Cartesian coordinates to cylinder coordinates. As shown in Fig. 5, during rasterization, a small fragment program transforms rasterization coordinates to cylindrical coordinates, performs a (bilinearly interpolated) look-up into the 2D emission density texture map, and accumulates the emission. To add up all rendered slices without introducing rounding errors due to limited precision, we also make use of the floating-point buffer available on modern graphics hardware.

The accuracy of the rendered image is determined by the interdependence between rendering resolution, the number of rendered slices, and emissive density map resolution. For convenient comparison, rendering resolution is matched to the resolution of the reference photograph. To minimize aliasing, the number of slices should be equivalent to the linear dimension of the rendered image. The emissive density map, in contrast, should exhibit half the resolution of the reference image, as described in Section 5. An nVidia GeForce FX 3000 graphics card easily attains 10 fps when rendering 128 depth layers at 128 $\times$ 128-pixel image resolution. Assuming 10^6 iterations, our CIVR algorithm converges in approximately one day.

4.2 Projection Error Evaluation

The aim of CIVR is to determine the emission density map entries such that its 2D projection, the volume rendering

```
1   MOV R0, f[TEX0];
2   MUL R1, R0,R0;
3   ADD R1.z, R1.x,R1.y;
4   RSQ R0.x, R1.z;
5   RCP R0.x, R0.x;
6   MAD R0.y, R0.z, 0.5, 0.5;
7   MOV R0.z, R0.w;
8   MUL R0,R0, p[0];
9   TEX R1, R0.xyzw, TEX0, RECT;
10  TEX R0, f[TEX1].xyzw, TEX1, RECT;
11  ADD o[COLH], R1,R0;
12  END;
```

Fig. 5. Source code of the CIVR fragment program. As input, texture unit 0 is assigned the emission density texture map. The unit's texture coordinates already correspond to the model's frame of reference, having been transformed by the appropriate texture matrix. Texture unit 1 represents the floating-point (P) buffer, its texture coordinates correspond to the fragment's position in the buffer. The constant `p[ 0]` stores the width and height of the emission density texture map in the x and y component, respectively. The program converts the fragment's Cartesian coordinates in the model frame to cylindrical coordinates with a radial and an axial component (`R0.x` and `RO.y` in lines, 5, 6). The coordinates are scaled to match the size of the emission density map (line 8) and the emission map is queried (line 9). Additionally, the fragment position's corresponding value in the floating-point buffer is read (line 10) before both values are summed and passed on as fragment color (line 11).

result, matches the true PN image as closely as possible. This is accomplished by comparing the rendering result to the original PN image and systematically varying the emission density map values until the difference is minimal.

To quantify the difference between rendered model and actual photograph, the floating-point buffer is read to main memory. Each rendered pixel is compared to its corresponding pixel in the PN image, and the sum-of-squared-differences over all pixels is computed. We have experimented with different error measures before returning to sum-of-squared-differences, which yielded the fastest convergence results. Since emission density cannot become negative, a very high error value is returned whenever the optimization routine attempts to set a map element to a negative value. This allows the otherwise generic optimization routine to be constrained to a physically plausible PN model.

4.3 Nonlinear Optimization

From an optimization-theoretical point of view, a particular emission density map in conjunction with the nebula's 3D orientation angles constitutes one point in the high-dimensional model parameter space. For initialization, all density map elements are set to zero, the orientation angle is assigned the value that is determined from the images, Section 5, and the inclination angle is set to $90°$, i.e., the symmetry axis is assumed to be orthogonal to the viewing direction. From this starting point, the optimization routine then traverses the model parameter space and converges toward the (local) minimum of the SSD error function.

Since we do not know the nebula's true 3D orientation in space, the relationship between the axisymmetric density map and the nebula's observed projection image is nonlinear. Also, map entries can only have nonnegative values. We therefore employ a nonlinear optimization approach that relies on repeatedly evaluating an error function measuring the difference between the rendered PN model and the photograph of its celestial counterpart. We employ a standard implementation of Powell's optimization method [34]. It numerically evaluates the error function's local slope along all dimensions from which it determines the conjugate-gradient direction. Each optimization iteration step entails modifying the density map elements and orientation angles, uploading the new map values to the graphics card, volume-rendering the model, and recalculating the error measure. This analysis-by-synthesis loop iterates until the algorithm converges to a minimum of the error function. It must be noted that, since we are dealing with a nonlinear optimization problem, convergence to the global minimum cannot be assured. Due to the model constraints, however, in our experiments with real as well as simulated data, Section 7, we observe that the model is always robustly driven toward a well-matching minimum.

4.4 Convergence Acceleration

Naively, the emission density map may be optimized by considering the full-resolution density map directly such that, per iteration step, one density map element is optimized individually. This approach, however, does not take into account the influence of other density map elements. Oscillating element values are the result and convergence is prohibitively slow. Instead, a multiresolution representation of the emission density map is employed which also has a regularizing effect on the optimization result.

In a first attempt, we represented the density map using Haar-wavelet basis functions. The wavelet coefficients were optimized in order of ascending spatial frequency. While convergence performance was much improved in comparison with the naive full-resolution approach, the wavelet coefficients had to be transformed back to the 2D map domain prior to each SSD error evaluation to enforce nonnegative values of all density map elements. In order to avoid this additional computational overhead in the innermost loop of the optimization routine, instead, we found that a simple resolution pyramid performs just as well. Fig. 6 depicts the hierarchical reconstruction approach. Reconstruction starts from a low-resolution version of the density map and the image. At this level, only a fraction of density map entries must be reconstructed and the model is volume-rendered very quickly at low resolution. After convergence, the reconstructed low-resolution density map is scaled up to serve as initialization for the next-higher resolution density map. When doubling the map's resolution, rendering resolution must be increased likewise, which in turn increases model rendering time. Employing four hierarchy levels, the reconstruction of a 128×32-pixel density map takes approximately one day on a 2.4 GHz PC in conjunction with an nVidia GeForce FX 3000 graphics card.

5 Reconstruction

For PN reconstruction, the CIVR algorithm is applied to photos of planetary nebulae recorded at optical wavelengths. Mathematically, the recovery of the axisymmetric emission density distribution from a 2D image is equivalent to deprojecting the image, i.e., a deconvolution

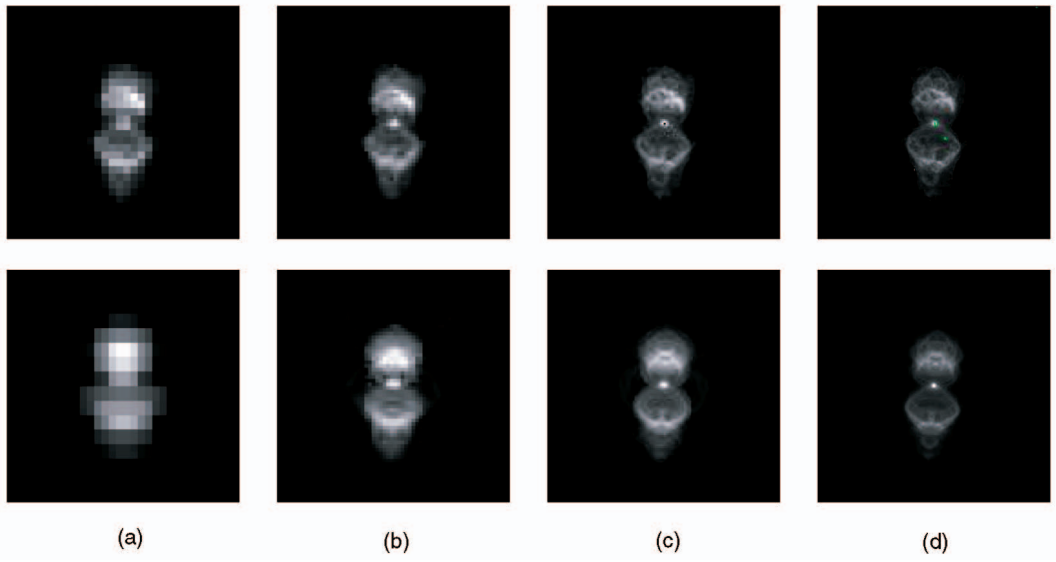

Fig. 6. Multiresolution optimization: The upper row depicts the actual image of the hydrogen emission line of Mz-3 at different resolution levels, the lower row shows the corresponding rendered model after convergence. The sampling theorem dictates that the model can be reconstructed only at half the resolution of the respective image. CIVR starts out with optimizing a low-resolution version of the emission density map (a). After convergence, the result is then up-sampled and used as an initial map for optimizing at the next-higher resolution level (b), (c), (d). The full-resolution version (d) closely matches the actual photograph of the planetary nebula (e). In the photo, the central star and background stars are masked (green) prior to running the CIVR algorithm.

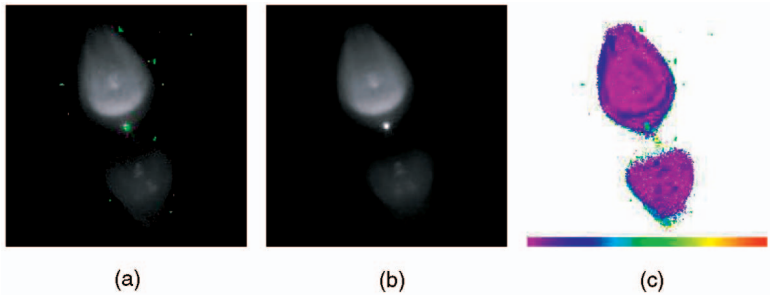

Fig. 7. Prior to applying CIVR, the central star and background stars in the original photograph (a) are masked out (green). The reconstructed axisymmetric nebula model (b) closely resembles the original image. The difference image (c) depicts local contrast between (a) and (b). Purple represents little contrast (difference) between (a) and (b), red denotes maximum contrast.

with a point spread function that depends on the (unknown) 3D orientation of the nebula. This deprojection becomes possible because of the nebula's axial symmetry as long as the symmetry axis is not pointed toward our viewing position. For axis-on PNe, the axisymmetric emission density distribution cannot be unambiguously recovered.

Constrained inverse volume rendering takes optical images of planetary nebulae as input. Research-grade images of many PNe are publicly available in FITS file format on the Internet [35]. The data used for our work, and the astronomical images shown here, are from the Hubble Space Telescope Archive [27] and the IAC Morphological Catalog of Northern Galactic Planetary Nebulae [28]. The images have been acquired with narrow-bandpass filters at distinct wavelengths that correspond to specific chemical elements. The scalar-valued images feature 16 bit-accurate linear dynamic range.

Prior to applying our CIVR volume modeling method, the images must be preprocessed to correct for a number of imaging-related effects in order to obtain scientifically accurate data. For our CIVR technique to produce correct results, the images must be calibrated to exhibit a linear response curve, i.e., twice the number of photons falling on a pixel must double the pixel's value. We calibrate the raw image data using the standard astronomical image processing software IRAF [36] and STSDAS [37]. Stars nearby and behind the nebula interfere with the reconstruction process because their light originates outside the nebula model. Such regions are manually masked. Fig. 7a depicts the preprocessed photograph of M1-92 taken in the light of nitrogen, with the masked regions highlighted in green. To avoid bias in the reconstruction result, the CIVR algorithm disregards these regions when computing the error function; the masked pixels are effectively missing from the image. Due to imaging noise, the background pixels in the images additionally exhibit small, nonzero values. This is a second potential source of inaccuracy since the CIVR algorithm would attempt to model the pixels outside the nebula as part of the nebula itself. Thus, we threshold the images and set the background uniformly to zero. Finally, the PN's central star is identified in each image and the nebula's projected axis of symmetry is approximately determined. This can be done automatically by finding

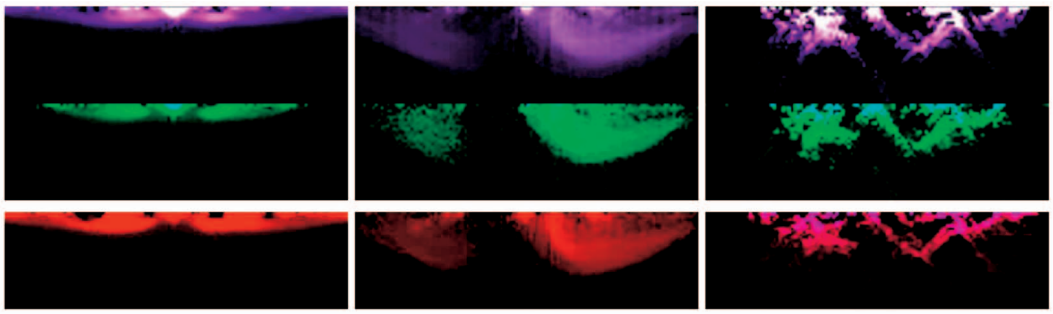

Fig. 8. Axisymmetric emission density maps for the planetary nebulae He2-437, M1-92, and Mz-3 (left to right). From top to bottom, the maps correspond to the distribution of hydrogen, oxygen, and nitrogen (sulfur for Mz-3). Empty voids are apparent within He2-437, while luminous wall-like structures dominate Mz-3. M1-92 is more homogeneous. These differences may be related to the central star's mass and the processes governing the gas outflow, as well as due to interactions with interstellar matter.

the maximum pixel value in the image (corresponding to the central star's position) and calculating the principal eigenvectors of the image's covariance matrix. The PN image is centered on the central star and clipped to 256×256 pixels, which is sufficiently large to completely accommodate the nebulae in all our test images.

PNe are small astronomical objects in the sky, typically measuring no more than a couple of arcminutes in diameter. Since atmospheric turbulence limits useful imaging resolution from ground-based telescopes to about one arcsecond, sensibly, PN images are no larger than a few hundred pixels along their longest dimension. To determine at which resolution our axisymmetric density map may be optimally reconstructed, we employ the sampling theorem and make sure that the highest measurable frequency corresponds to two samples per period. It follows that, for a nebula covering 256 pixels in the image along its longest direction, we should not attempt to recover more than 128 elements in our density map.

We now apply our CIVR algorithm to each calibrated gray-scale photograph individually. Initially, all density map entries are set to zero. Because the nebula's symmetry axis goes through the central star, we can initialize the orientation angle based on the slope of the PN's line of symmetry in the image. The second, unknown angle, representing the inclination of the nebula's symmetry axis toward Earth, is initialized to be perpendicular to our viewing direction. During CIVR, both angles are optimized to determine the true 3D orientation of the nebula, in addition to the emission density map entries. Fig. 7b shows the final nebula model rendered from the same viewing direction as the original photograph, Fig. 7a. Note that the masked (green) regions in Fig. 7a have not been taken into account during optimization, so these regions are not necessarily expected to match. To illustrate the residual differences, Fig. 7c depicts local contrast between the original image and the model rendering. Local contrast c at pixel coordinates (x, y) is defined as

$$c(x,y) = \frac{\|p_I(x,y) - p_R(x,y)\|}{p_I(x,y) + p_R(x,y)}, \qquad (1)$$

where p_I is the pixel intensity value in the original image and p_M is the corresponding rendered pixel value. Note that p_I and p_M are always nonnegative. When $c = 0$, contrast vanishes, signifying an exact match between both pixel values, whereas maximum contrast $c = 1$ denotes maximum deviation. In Fig. 7c, local contrast c is linearly color-encoded, with purple denoting small contrast values and red representing maximum contrast $c = 1$. As can be seen, the rendered model matches the photograph well over most of the nebula's area. Deviations along the edges of the dim lower gas lobe can be attributed to image noise. The small differences in the upper lobe showing up in blue, on the other hand, may be a sign of absorption due to localized dust clouds, breaking the nebula's axisymmetry.

6 Results

We evaluate our CIVR algorithm on photographs of three different PNe. For each planetary nebula, several images are used, taken at different wavelengths and corresponding to different chemical elements, as shown in Fig. 1. Using CIVR, for each element, a separate axisymmetric density map is reconstructed. Fig. 8 depicts each element's spatial distribution within the nebulae. The three-dimensional gas distribution is obtained by revolving each density map about its upper edge.

In addition to the emission density maps, the orientation angles of each nebula are determined as stated in Table 1. Note that orthographic projection implies that the inclination angle can only be recovered up to a two-fold ambiguity, yielding symmetrical solutions with respect to the perpendicular orientation. If we assume an equatorial dust belt around the central star, however, this symmetry can be broken as the dimmer of either lobe must be behind the absorbing dust. This effect might be responsible for the difference in brightness of the lobes of M1-92 in Fig. 9.

TABLE 1
Reconstructed Nebula Oritentations

PN	incl. angle	orient. angle
He2-437	84°/96°	77.5°
M1-92	79°/101°	16.5°
Mz-3	55°/125°	175.5°

The inclination angle can only be determined up to a symmetry ambiguity with respect to the perpendicular $90°$ orientation.

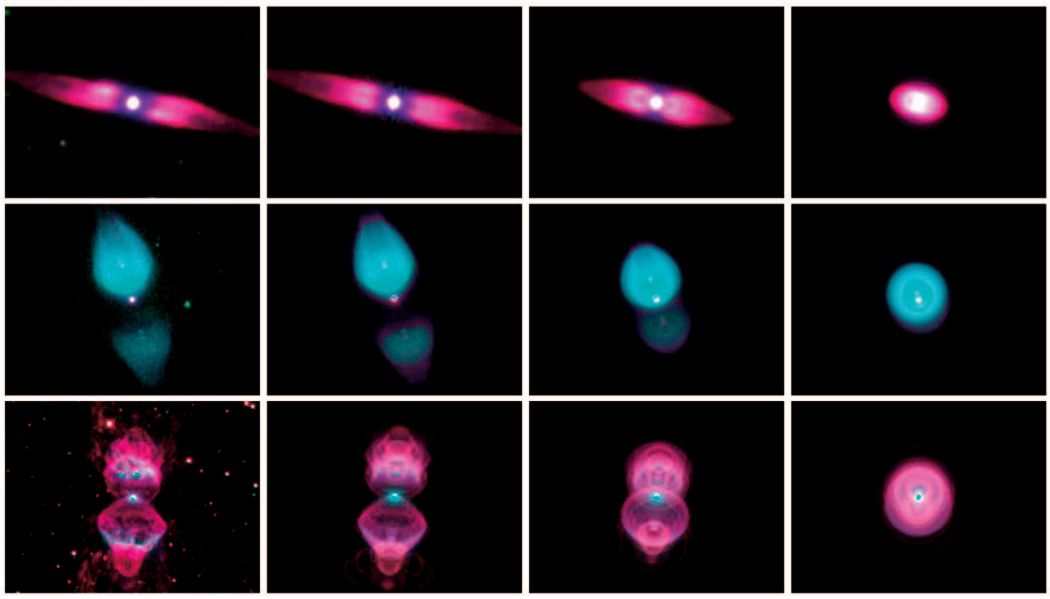

Fig. 9. CIVR results for the planetary nebulae He2-437 (top), M1-92 (middle), and Mz-3 (bottom). Left to right: The color-composited original image, the same view rendered from the reconstructed axisymmetric model, and two "unearthly" sights of the nebulae from inclination angles 35° and 10°.

6.1 Realistic 3D Visualization

Our reconstructed PN models enable us to accurately visualize these attractive objects from arbitrary viewpoints, useful, e.g., for planetarium shows or interactive educational tools. Fig. 9 depicts color composites of the actual PN images along with renderings from the same viewpoint obtained from the reconstructed models as well as visualizations of the nebulae as they would appear from vantage points thousands of light years from Earth.

To achieve the most natural visual impression, we made an effort to reproduce the nebula's actual colors. For each PN, we have available images corresponding to one characteristic emission line for the elements hydrogen, oxygen, and nitrogen or sulfur, shown in Fig. 1. These are the elements responsible for the visible color appearance of PNe. We assume that the images have been obtained using narrow-band filters which single out only the respective element's emission line. While relative emission line strengths are somewhat dependent on the physical conditions inside the nebula (temperature, electron density), we have adopted here mean values for PNe as published in the literature [24], [25]. Having rendered a view using all elements' emission density maps, for each pixel, we construct its spectrum according to the element emissions' intensities [5], which we convert to RGB values using the CIE 1931 XYZ color matching function values and the appropriate color space transformation [38].

Planetary nebula images exhibit a wide dynamic range. While the central star typically represents the brightest region in any PN image, the outermost regions of glowing gas frequently drown in detector noise. To simultaneously display the bright regions and the faint structure of the gas, we employ simple gamma correction to all three color channels. More advanced tone mapping operators may yield even more visually attractive rendering results [39]. The right columns in Fig. 9 depict our test PNe as they would appear under different inclination angles, corresponding to viewpoints located thousands of light years from Earth [40]. Notice the considerable change in appearance with orientation angle, but also the diversity among different planetary nebulae. By studying the variety in possible PN appearance, the CIVR results may help astronomers identify planetary nebulae seen at low inclination angles that have so far eluded identification.

6.2 Astrophysical Analysis

An astrophysical interpretation of the reconstructed axisymmetric maps must be done with care. Our neglect of absorption in and around the nebula can especially lead to systematic errors. This is obvious for M1-92 in Fig. 9, whose lower lobe appears much dimmer than the upper lobe, which is probably due to an equatorial belt of absorbing dust around the central star. On the other hand, the accuracy of our fitted model, depicted in Fig. 7c, is a strong indication that any dust present must be located outside the emissive volume and obey the nebula's axisymmetry, in which case our reconstruction method remains applicable.

Visualizing the difference between the original image and the reconstructed model of He2-437, deviations from the axisymmetry assumption become apparent in Fig. 11. The aberrations from axial symmetry show an orderly pattern, resembling higher-order harmonics. This can be taken as an indication that fundamental physical processes are the cause of the anomalies, for example, a precessing central star.

Gas emission density varies considerably within each nebula, indicating that the gas is rather inhomogeneously distributed, as can be observed in Fig. 8. The internal structure differs considerably among the nebulae: Large empty voids centered on the symmetry axis are present in He2-437. M1-92 is more homogeneous, featuring a smooth density gradient toward the luminous region close to the central star in the right lobe. The distribution in Mz-3 is even more complex,

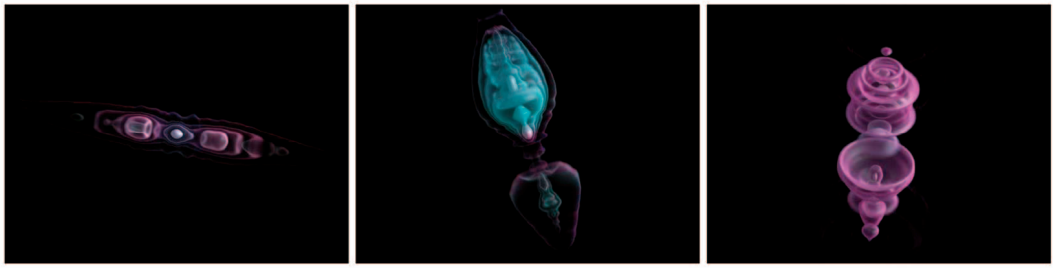

Fig. 10. Volume renderings to visualize 3D isosurfaces in He2-437 (left), M1-92 (middle), and Mz-3 (right). The renderings make apparent the spatial inhomogeneities and internal structure of the nebulae.

with most emission being concentrated in thin veils, probably along the shock front. To help convey the three-dimensional structure of the PNe, Fig. 10 depicts volume renderings of the nebulae using multiple soft isosurfaces of the emission density. The coloring on the isophotes originates from a volume RGB texture generated similarly to the images in Fig. 9 Diffuse shading aids in interpreting 3D shape, making apparent the spatial inhomogeneities. This can be seen for M1-92 whose seemingly homogeneous internal structure in Fig. 8 turns out to be much more complex. Several nested shells and emission concentrations along the axis of symmetry become apparent in the volume visualization. For He2-437, the visualization shows that the empty bubbles within the nebula are surrounded by strong emission, probably gas concentrations that are caused by magnetic fields which are also responsible for the nebula's constricted shape. Mz-3, finally, shows open surfaces of glowing gas along the shock front that may be due to an eruptive outbreak of Mz-3's central star.

To investigate the spatial correlation of different chemical elements within the nebulae, Fig. 12 illustrates the deviation of relative element abundance for oxygen and nitrogen (sulfur for Mz-3) with respect to hydrogen. The reconstructed density maps for the three nebulae in Fig. 8 are used to compute the relative element abundance

$$a_{O,N,S/H}(l,\rho) = \frac{m_{O,N,S}(l,\rho)}{m_H(l,\rho)} - \bar{a}_{O,N,S/H},$$

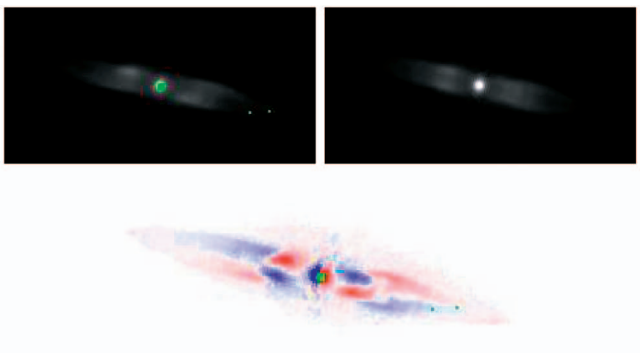

Fig. 11. Deviation from axisymmetry. Top: He2-437 in the light of hydrogen ($H\alpha$) and the corresponding rendered axisymmetric model. Bottom: Color-encoded difference image. Negative differences are depicted in blue, positive difference in red. The symmetric deviation from axial symmetry may be a sign of the central star's precessing magnetic field.

where $m_{O,N,S}$ is the value of the density map for oxygen, nitrogen or sulfur at position (l,ρ), m_H is the value in the hydrogen map, and $\bar{a}_{O,N,S/H}$ is the average relative abundance over all map elements exceeding a minimum threshold. Fig. 12 depicts regions of smaller-than-average relative abundance in blue and regions of larger-than-average relative concentration in red. Spatial correlation differs substantially between nebulae as well as between elements: In He2-437, relative abundance of hydrogen and oxygen is almost homogeneous, while nitrogen and hydrogen are distributed distinctly differently. In M1-92, on the other hand, all three elements have similar relative abundances within the entire nebula. In contrast, relative element concentrations in Mz-3 vary very strongly over space. Astrophysical models attempting to describe the formation and evolution of planetary nebulae must be able to explain these differences.

7 Validation and Discussion

The CIVR method reconstructs the axisymmetric structure of PNe such that the reprojection is a very close match to the actual, recorded image. In the following, we validate the credibility of CIVR results and discuss the inherent limitations of our reconstruction approach.

Since our reconstruction problem is nonlinear, a proof for general convergence to the global minimum of our error function cannot be given. To validate the reconstruction results derived for real nebulae, we would have to know the actual 3D distribution. While there exists a rather involved, indirect observation technique known as echellogrametry [18], [19] to reconstruct 3D gas velocities in spherical nebulae, no measured 3D gas distribution data is available for planetary nebulae with a pronounced axisymmetry.

Instead, we must rely on synthetic images to test our approach. We make use of Aaquist and Kwok's ellipsoid shell model [32] to generate test images for a number of different orientations. We then let our CIVR algorithm recover the axisymmetric density distribution and orientation angles and compare the reconstructed values with the original values derived from the ellipsoid shell model.

Fig. 13a shows a synthetic density map with a corresponding projection image. The CIVR-reconstructed density map and rendered model image are depicted in Fig. 13b. While the original density map is not recovered exactly, Fig. 13c (top), the close resemblance shows that the found values are in the vicinity of the exact solution. The largest

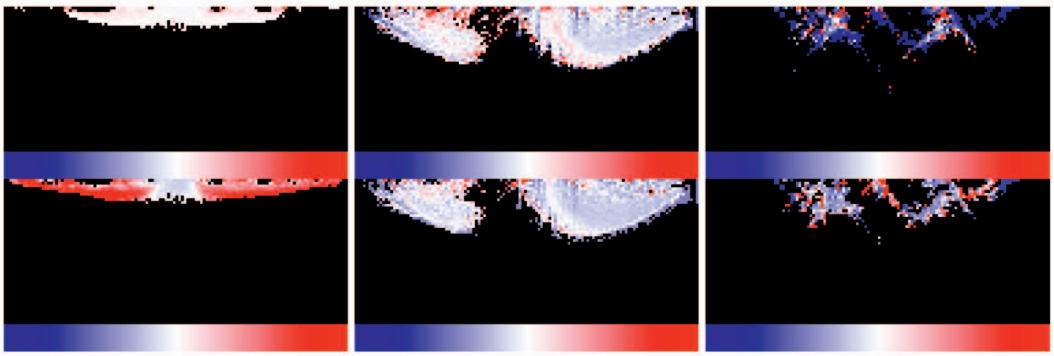

Fig. 12. Spatial correlation of chemical elements: The figures illustrate the relative amount of oxygen (upper row) and nitrogen/sulfur (lower row) with reference to hydrogen in the axisymmetric density maps of He2-437, M1-92, and Mz-3 (left to right). Deviations from the average percentage are depicted in blue or red, corresponding to smaller-than-average or large-than-average local element abundance.

errors occuring along the upper and lower map boundaries are due to residual aliasing.

For inclination angles from $90°$ down to $40°$, the reconstructed map values are within the error range that can be expected from the optimization routine's preset termination error threshold. The accuracy of the recovered inclination angle varies with angle value, from about $3°$ error for $90°$ inclination down to less than one degree at $45°$ inclination. While we cannot test all possible cases, our validation experiments indicate that CIVR reconstruction results do, indeed, closely resemble the true 3D structure of planetary nebulae.

CIVR can be applied not only to axisymmetric PNe, but also to spherical nebulae by exploiting their point symmetry. Our reconstruction approach is not applicable for irregularly shaped nebulae, for example, if the outflowing gas collides with interstellar gas clouds. The method works best for nebulae seen at high inclination angles, i.e., for PNe whose axis of symmetry is oriented approximately perpendicular to our viewing direction. This can be understood by approximating the reconstruction as an inverse linear problem and looking at the condition number of the projection matrix. With decreasing inclination angle, i.e., when the PN's symmetry axis approaches the viewing direction, the condition number increases and the (unregularized) reconstruction result becomes unstable. At $0°$ inclination, the symmetry axis points directly toward the observer and the nebula's axial symmetry is of no use to constrain the reconstruction problem: The linear system of equations is underdetermined. Our validation experiments indicate that CIVR is capable of robustly reconstructing PNe down to inclination angles of $40°$.

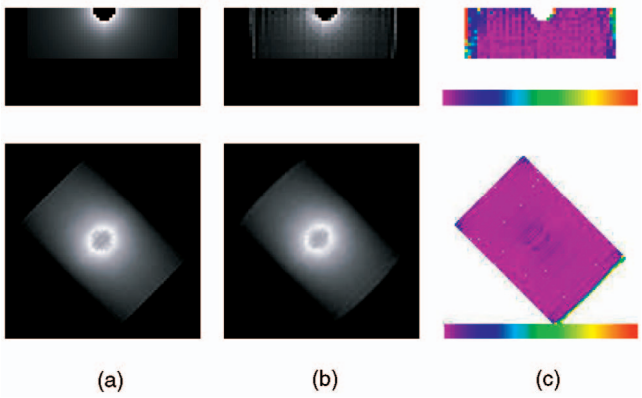

Fig. 13. CIVR validation on synthetic data. Using the ellipsoid shell model [32], we generate the axisymmetric density map of a synthetic planetary nebula (a, top) and render its appearance for $0°$ inclination, $45°$ orientation (a, bottom). From the projection image (a, bottom), the density map is reconstructed using CIVR (b, top) such that the rendered model (b, bottom) closely resembles the input image (a, bottom). The difference between the ground-truth data (a) and the reconstruction (b) is depicted in (c). Shown is the color-coded contrast (1). Purple denotes minimal difference between (a) and (b), while red represents maximum contrast.

8 Conclusions

We have presented a constrained inverse volume rendering (CIVR) approach to recovering the three-dimensional distribution of the photo-ionized gas in planetary nebulae (PNe) from optical images. Their inherent axial symmetry enables us to use an axisymmetric model to describe the nebula's emission distribution in space, thereby constraining the reconstruction problem. Being purely emissive, the PN model is volume-rendered efficiently using modern PC graphics hardware. By comparing the rendering result with the actual image, an optimization algorithm varies local emission as well as orientation angles to match the rendered model to the recorded image.

From the reconstructed spatial emission distribution, the nebulae can be rendered from arbitrary perspective. Applications include creating realistic visualizations of PNe as illustration material for astronomy courses, science programs, planetarium shows, or TV documentaries. Collaborations with a number of planetariums are underway. Furthermore, the recovered axisymmetric shape is useful for studying the spatial relationships of element abundances, as well as to validate and refine the interacting stellar winds theory. By comparing CIVR reconstruction results to photoionization and gas-dynamic simulations [22], astrophysical models of planetary nebulae can be validated. This might be of help to establish a general evolutionary sequence among observed PNe that can explain the variations in PN composition and appearance based on only a handful of physical properties of the central star.

ACKNOWLEDGMENTS

The authors are indebted to Sun Kwok and his student Ting-Hui Lee at the University of Calgary, Canada, to Stefan Kimeswenger at the University of Innsbruck, Austria, as well as to Franco Sabbadin at Padova Observatory, Italy, for their very helpful comments and their constructive feedback on the usefulness of the presented method for astrophysics research. Their thanks also go to Bastian Goldlücke for his help in programming around the graphics hardware driver's bugs and to Andrei Lintu for calibrating the astronomical images. This project was funded in part by the German Science Foundation DFG, contract no. 444USA111/2/04.

REFERENCES

[1] H. Curtis, "The Planetary Nebulae," *Publ. Lick Observatory, part III,* no. 13, pp. 57-74, 1918.
[2] G. Khromov and L. Kohoutek, "Morphological Study of Planetary Nebulae," *Planetary Nebulae,* D. Osterbrock and C. O'Dell, eds., pp. 227-235, 1968.
[3] B. Balick, "The Evolution of Planetary Nebulae. I—Structures, Ionizations, and Morphological Sequences," *Astronomical J.,* vol. 94, pp. 671-678, Sept. 1987.
[4] C. Masson, "The Structure of and Distances to BD + 30 deg 3639 and NGC 6572," *Astrophysical J.,* vol. 346, pp. 243-250, 1989.
[5] M. Magnor, G. Kindlmann, C. Hansen, and N. Duric, "Constrained Inverse Volume Rendering for Planetary Nebulae," *Proc. IEEE Visualization 2004,* pp. 83-90, Oct. 2004.
[6] B. Curless and M. Levoy, "A Volumetric Method for Building Complex Models from Range Images," *Proc. ACM SIGGRAPH (Computer Graphics Series),* vol. 30, pp. 303-312, 1996.
[7] K. Mueller, R. Yagel, and J. Wheller, "A Fast and Accurate Projection Algorithm for the Algebraic Reconstruction Technique (ART)," *Proc. SPIE Medical Imaging Conf.,* vol. SPIE 3336, pp. 724-732, Feb. 1998.
[8] S. Marschner, "Inverse Rendering for Computer Graphics," PhD dissertation, Cornell Univ., 1998.
[9] M. Levoy, "Display of Surfaces from Volume Data," *IEEE Computer Graphics and Applications,* vol. 8, no. 5, pp. 29-37, 1988.
[10] B. Cabral, N. Cam, and J. Foran, "Accelerated Volume Rendering and Tomographic Reconstruction Using Texture Mapping Hardware," *Proc. Symp. Volume Visualization (VolVis '94),* pp. 91-98, 1994.
[11] C. Rezk-Salama, K. Engel, M. Bauer, G. Greiner, and T. Ertl, "Interactive Volume Rendering on Standard PC Graphics Hardware Using Multi-Textures and Multi-Stage Rasterization," *Proc. ACM SIGGRAPH/EUROGRAPHICS Workshop Graphics Hardware,* pp. 109-118, 2000.
[12] N. Max, "Optical Models for Direct Volume Rendering," *IEEE Trans. Visualization and Computer Graphics,* vol. 1, no. 2, pp. 99-108, 1995.
[13] D. Nadeau, J. Genetti, S. Napear, B. Pailthorpe, C. Emmart, E. Wesselak, and D. Davidson, "Visualizing Stars and Emission Nebulae," *Computer Graphics Forum,* vol. 20, no. 1, pp. 27-33, Mar. 2001.
[14] J. Genetti, "Volume-Rendered Galactic Animations," *Comm. ACM,* vol. 45, no. 11, pp. 62-66, Nov. 2002.
[15] W. Zheng and C. O'Dell, "A Three-Dimensional Model of the Orion Nebula," *Astrophysical J.,* vol. 438, no. 2, pp. 784-793, Jan. 1995.
[16] S. Kwok, C. Purton, and P. Fitzgerald, "On the Origin of Planetary Nebulae," *Astrophysical J.,* vol. 219, pp. L125-L127, 1978.
[17] N. Calvet and M. Peimbert, "Bipolar Nebulae and Type I Planetary Nebulae," *Revista Mexicana de Astronomia y Astrofisica,* vol. 5, p. 319, 1983.
[18] F. Sabbadin, "Spatiokinematic Models of Five Planetary Nebulae," *Monthly Notices of the Royal Astronomical Soc.,* vol. 210, pp. 341-358, 1984.
[19] F. Sabbadin, E. Cappellaro, S. Benetti, M. Turatto, and C. Zanin, "Tomography of the Low Excitation Planetary Nebula NGC 40," *Astronomy & Astrophysics,* vol. 355, pp. 688-698, Mar. 2000.
[20] W. Saurer, "Morphology and Expansion Characteristics of the Planetary Nebula M 1-79," *Astronomy & Astrophysics,* vol. 326, pp. 1187-1194, Oct. 1997.
[21] C. Mellema and A. Frank, "Radiation Gasdynamics of Planetary Nebulae—V. Hot Bubble and Slow Wind Dynamics," *Monthly Notices of the Royal Astronomical Soc.,* vol. 273, pp. 401-410, 1989.
[22] B. Armsdorfer, S. Kimeswenger, and T. Rauch, "Effects of Central Star Models on Planetary Nebulae Shell Modelling," *Ionized Gaseous Nebulae,* Nov. 2000.
[23] C. Ercolano, C. Morisset, M. Barlow, P. Storey, and X.-W. Liu, "Three-Dimensional Photoionization Modelling of the Planetary Nebula NGC 3918," *Monthly Notices of the Royal Astronomical Soc.,* vol. 340, pp. 1153-1172, 2003.
[24] D. Osterbrock, *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei.* Sausalito, Calif.: University Science Books, 1989.
[25] S. Kwok, *The Origin and Evolution of Planetary Nebulae.* Cambridge, U.K.: Cambridge Univ. Press, 2000.
[26] B. Balick and A. Frank, "Shapes and Shaping of Planetary Nebulae," *Ann. Rev. Astronomy and Astrophysics,* vol. 40, pp. 439-487, 2002.
[27] Hubble, 2004, http://archive.stsci.edu/hst/search.php.
[28] IAC, 2004, http://www.iac.es/nebu/nebula.html.
[29] A. Frank, B. Balick, V. Icke, and G. Mellema, "Astrophysical Gasdynamics Confronts Reality—The Shaping of Planetary Nebulae," *Astrophysical J.,* vol. 404, no. 1, pp. L25-L27, 1993.
[30] F. Kahn and K. West, "Shapes of Planetary Nebulae," *Monthly Notices of the Royal Astronomical Soc.,* vol. 212, pp. 837-850, Feb. 1985.
[31] C. Morrisset, R. Gruenwald, and S. Viegas, "Morphology and Kinematics of Planetary Nebulae. I A New Modeling Tool," *Astrophysical J.,* vol. 533, pp. 931-937, Apr. 2000.
[32] O. Aaquist and S. Kwok, "Radio Morphologies of Planetary Nebulae," *Astrophysical J.,* vol. 462, no. 2, pp. 813-824, May 1996.
[33] C. Zhang and S. Kwok, "A Morphology Study of Planetary Nebulae," *Astrophysical J. Supplement Series,* vol. 117, no. 2, pp. 341-359, Aug. 1998.
[34] W. Press, B. Flannery, S. Teukolsky, and W. Vetterling, *Numerical Recipes in C.* Cambridge Univ. Press, 1992.
[35] NASA/GSFC, "Flexible Image Transport System (FITS) Support Office,"http://fits.gsfc.nasa.gov/, 2004.
[36] IRAF, 2004, http://iraf.tuc.noao.edu.
[37] STSDAS, 2004, www.stsci.edu/resources/software_hardware/stsdas.
[38] G. Wyszecki and W.S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae,* second ed. New York: John Wiley & Sons, 1982.
[39] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, "Photographic Tone Reproduction for Digital Images," *ACM Trans. Graphics,* vol. 21, no. 3, pp. 267-276, 2002.
[40] M. Magnor, "CIVR Demo Videos," http://www.mpi-sb.mpg.de/magnor/civr.html, 2004.

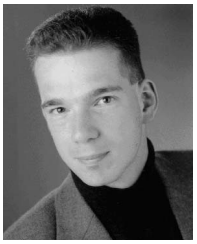

Marcus Magnor received the BS in 1995 from the University of Würzburg, Germany, the MSc degree in physics from the University of New Mexico in 1997, and the PhD degree in electrical engineering from the University of Erlangen, Germany, in 2000. He is head of the independent research group "Graphics-Optics-Vision" at the Max-Planck-Institut Informatik in Saarbrücken, Germany. Before his current position, he spent 2001 as a research associate in the Computer Graphics Lab at Stanford University, Stanford, California. His research interests include video-based rendering and inverse rendering approaches, as well as data-based modeling techniques. He is a member of the IEEE.

Gordon Kindlmann received the BA degree in mathematics and the MS degree in computer graphics from Cornell University in 1995 and 1999, respectively, and the PhD degree in computer science, under Christopher R. Johnson, from the University of Utah in 2004. He is currently a postdoctoral research fellow in the Laboratory of Mathematics in Imaging at Brigham and Women's Hospital in Boston. His research focuses on the visualization and analysis of diffusion tensor fields, with continuing work on transfer functions and volume rendering, and perceptual issues in scientific visualization.

Charles Hansen received the BS degree in computer science from Memphis State University in 1981 and the PhD degree in computer science from the University of Utah in 1987. He is an associate professor of computer science at the University of Utah. From 1997 to 1999, he was a research associate professor in computer science at Utah. From 1989 to 1997, he was a technical staff member in the Advanced Computing Laboratory (ACL) located at Los Alamos National Laboratory, where he formed and directed the visualization efforts in the ACL. He was a Bourse de Chateaubriand Postdoctoral Fellow at INRIA, Rocquencourt, France, in 1987 and 1988. His research interests include large-scale scientific visualization and computer graphics. He is a member of the IEEE.

Neb Duric has wide-ranging interests in imaging science. As an astrophysicist, he has published more than 100 articles and authored several books. His most recent book, *Advanced Astrophysics*, was published this year by Cambridge University Press. He recently joined Wayne State University and the Karmanos Cancer Institute in Detroit, Michigan, where he is a professor of radiation oncology and leads the development of a cancer imaging program. His current research activities include the development of ultrasound tomography and the reconstruction of images from tomographic data in support of early breast cancer detection.