

# Jupiter's shrinking Great Red Spot and steady Oval BA: Velocity measurements with the 'Advection Corrected Correlation Image Velocimetry' automated cloud-tracking method

Xylar S. Asay-Davis<sup>a,\*</sup>, Philip S. Marcus<sup>b</sup>, Michael H. Wong<sup>a,c</sup>, Imke de Pater<sup>a</sup>

<sup>a</sup> Department of Astronomy, University of California, Berkeley, CA 94720, USA

<sup>b</sup> Department of Mechanical Engineering, University of California, Berkeley, CA 94720, USA

<sup>c</sup> Space Telescope Science Institute, Baltimore, MD 21210, USA

## ARTICLE INFO

### Article history:

Received 31 March 2008

Revised 14 April 2009

Accepted 21 April 2009

Available online 18 May 2009

### Keywords:

Jupiter, atmosphere  
Atmospheres, dynamics  
Atmospheres, evolution

## ABSTRACT

We show that between 1996 and 2006, the area circumscribed by the high-speed collar of the Great Red Spot (GRS) shrunk by 15%, while the peak velocities within its collar remained constant. This shrinkage indicates a dynamical change in the GRS because the region circumscribed by the collar is nearly coincident with the location of the potential vorticity anomaly of the GRS. It was previously observed that the area of the clouds associated with the GRS has been shrinking. However, the cloud cover of the GRS is not coincident with the location of its potential vorticity anomaly or any other of its known dynamical features. We show that the peak velocities of the Oval BA were nearly the same in 2000, when the Oval was white, and in 2006, when it was red, as were all of the other features of the two velocities fields. To measure temporal changes in the GRS and Oval, we extracted velocities from images taken with Galileo, Cassini, and the Hubble Space Telescope using a new iterative method called Advection Corrected Correlation Image Velocimetry (ACCIV). ACCIV finds correlations over image pairs with 10-h time separations when other automated velocity-extraction methods are limited to time separations of 2 h or less. Typically, ACCIV velocities produced from images separated by 10 h had errors that are 3–6 times smaller than similar velocities extracted from images separated by 2 h or less. ACCIV produces velocity fields containing hundreds of thousands of *independent* correlation vectors (tie-points). Dense velocity fields are needed to locate the loci of peak velocities and other features.

© 2009 Elsevier Inc. All rights reserved.

## 1. Introduction

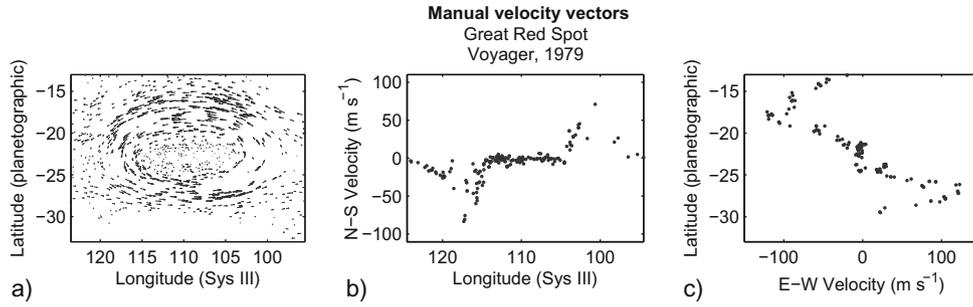
The Jovian atmosphere at the cloud-deck level has been changing. Between 1996 and 2006, the area of the cloud cover over the Great Red Spot (GRS) has been shrinking (Simon-Miller et al., 2002), and some observers suggest that the circumferential velocities of the Oval BA are increasing (Simon-Miller et al., 2006; Cheng et al., 2008). To verify or refute these and other changes and to analyze the uncertainties of the measurements of these changes, it is necessary to produce fields with large numbers of velocity vectors. This can only be done with automated methods. To illustrate the need for large numbers of velocity vectors, consider Fig. 1a, which shows all 1264 of the hand-derived velocity vectors of the GRS obtained from Voyager 1 by Sada et al. (1996). Although the number of vectors may seem large, Fig. 1b and c shows that this number is insufficient to deduce the locations of the peak velocities, the values of those peak velocities, or the spatial derivatives of the

\* Corresponding author. Address: Department of Astronomy, University of California – Berkeley, 601 Campbell Hall, Berkeley, CA 94720-3411, USA.

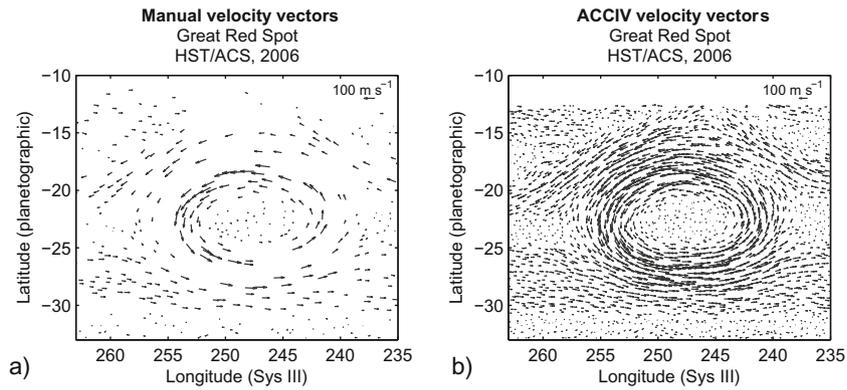
E-mail address: [xylar@astro.berkeley.edu](mailto:xylar@astro.berkeley.edu) (X.S. Asay-Davis).

velocities (i.e., the shear and vorticity), all of which are important for validating theoretical models and numerical simulations. To check that the sparsity of vectors in Fig. 1 was not due to the quality of the images, we extracted velocities manually and with our new automated method *Advection Correction Correlation Image Velocimetry* (ACCIV) using the *same* Hubble Space Telescope (HST) image pair with each method, and we obtained 443 vectors and ~140,000 independent vectors, respectively. (We defer our argument that the vectors are independent to Section 3.) The two extracted velocity fields are shown in Fig. 2 for comparison. Fig. 3 shows the manually-extracted velocities along the principal axes of the GRS using the vectors in Fig. 2a, which like the manually-extracted velocities in Fig. 1, are too sparse to obtain quantitative information. Thus, the sparsity of vector fields found with manual methods is systematic and problematic. On the other hand, the locations of the peaks and the derivatives of the velocities on the principal axes of the GRS can be determined accurately from the velocities extracted with ACCIV (Fig. 4).

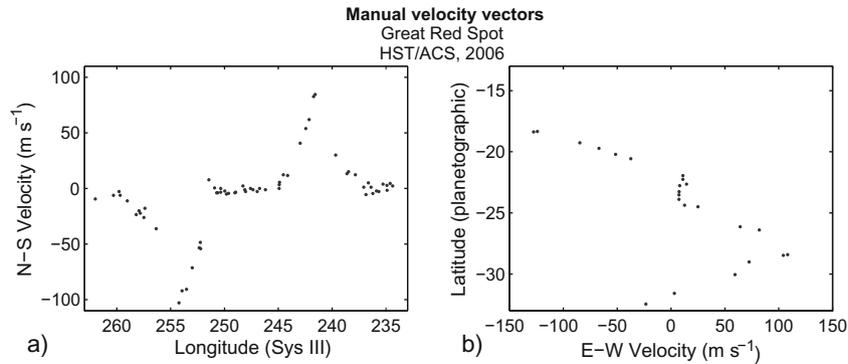
In most previous observational studies of Jovian vortices, velocity vectors were extracted from image pairs manually by visually identifying tie-points (or cloud displacement vectors), typically



**Fig. 1.** (a) All 1264 manual velocity vectors extracted by Sada et al. (1996) from the *VG1\_GRS\_79* data set (see Table 2). The  $1\sigma$  uncertainty (described in Section 3.6) of the data set is  $\sim 10 \text{ m s}^{-1}$ . (b) The north–south components of all velocity vectors within  $1^\circ$  of the east–west axis ( $23^\circ\text{S}$  planetographic latitude) of the GRS. (c) The east–west components of all velocity vectors within  $1^\circ$  of the north–south axis ( $109.5^\circ$  longitude) of the GRS. Using these figures, it is difficult to identify accurately dynamical characteristics such as the magnitudes and locations of the velocity peaks along the axes.



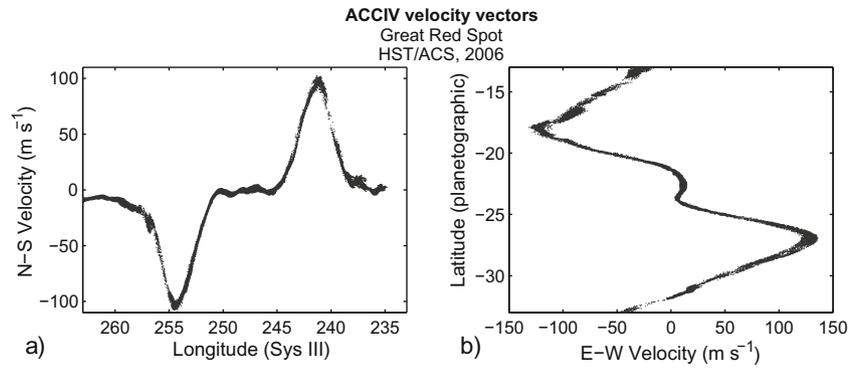
**Fig. 2.** (a) All 443 velocity vectors derived manually from the *HST\_GRS\_06* data set using images separated by about 10 h. We found the velocity vectors by tracking cloud features manually using the Matlab software described in Section 4.1. (b) Velocity vectors derived by ACCIV from the same images from the *HST\_GRS\_06* data set. For clarity, we have displayed  $\sim 3000$  out of the  $\sim 140,000$  independent velocity vectors in the complete velocity field.



**Fig. 3.** Principal axis velocities from the manually-derived velocity field shown in Fig. 2a. (a) The north–south components of all velocity vectors within  $1^\circ$  of the east–west axis ( $23^\circ\text{S}$  planetographic latitude) of the GRS. (b) The east–west components of all velocity vectors within  $1^\circ$  of the north–south axis ( $248^\circ$  longitude) of the GRS. The correlation uncertainty (described in Section 3.1) of the data set is  $\sim 8 \text{ m s}^{-1}$ .

yielding 100 to 1000 velocity vectors for a Jovian vortex (Table 1). Hand-derived velocity vectors are obtained by determining the displacement of a cloud feature over a time interval, assuming that the feature is a passive scalar advected by the ambient winds, and computing the velocity by dividing the displacement by the time interval. With hand-derived velocity vectors there is the satisfaction of knowing that there is a feature, identified by eye, associated with each velocity vector, but hand-derived velocity fields pose many problems in addition to their sparsity. Sub-pixel accuracy in finding the displacement vectors cannot be obtained, unless the hand-derived method is coupled to an automated correlation finder. Although there are publicly available programs for the tan-

dem use of hand and automated methods to obtain sub-pixel accuracy, such as AMOS (Yagi et al., 1978), such tandem use is generally not done (Simon-Miller, A., personal communication; Hueso, R. and Sánchez-Lavega, A., personal communication). Unless the time interval of the displacement is exceedingly small, or the path that the cloud feature travels is straight, rather than curved, the location assigned to a velocity vector that is derived by hand has large errors. Typically, the location assigned to a manually-derived velocity vector is half-way between the head and tail of the displacement vector. For curved paths, this is incorrect, and the correct location can only be determined by approximating spatial derivatives of the velocity field, which requires an interpolation



**Fig. 4.** Principal axis velocities from ACCIV applied to the *HST\_GRS\_06* data set (see Table 2 and Appendix B). The velocity vectors shown here and those in Fig. 2b are both subsets of the same ACCIV-extracted velocity field. (a) The north–south components of all velocity vectors within  $0.2^\circ$  of the east–west axis ( $23^\circ\text{S}$  planetographic latitude) of the GRS. (b) The east–west components of all velocity vectors within  $0.2^\circ$  of the north–south axis ( $248^\circ$  longitude) of the GRS. The correlation uncertainty of the data set is  $\sim 5\text{ m s}^{-1}$ .

**Table 1**

A summary of previous velocity-extraction methods. The manual methods were used to produce velocity fields of major Jovian vortices. The automated methods (distinguished as either 1-pass or 2-pass methods) have been applied to Jovian features including vortices, or to large areas of clouds on Jupiter, Saturn or Venus. Where applicable, data sets are described in more detail in Table 2. The separation time is the approximate time interval between images in the data set used to extract velocities. In most cases, the number of reported vectors is the number of velocity vectors in the vector field. (For Tokumaru et al. (1995), we list their reported number of degrees of freedom, which are not actually velocity vectors – see Section 2.) In some cases, the number of vectors was not reported, and we approximate (indicated by two or fewer significant digits) the number. We have only listed uncertainties that are based on the velocity data themselves and not simply on the properties of the image data set (see Section 7).

Author	Method	Data set ID or description	Sep. time (h)	Number of reported vectors	Mean vec. density ( $10^{-6}\text{ km}^{-2}$ )	Uncertainty ( $\text{m s}^{-1}$ )
Mitchell et al. (1981)	Manual	VG1_GRS_79	0.55	650	4	N/A
Mitchell et al. (1981)	Manual	VG1 Oval BC	1.39	350	3	N/A
Dowling and Ingersoll (1988)	Manual	VG1_GRS_79	10	2000	2	N/A
Dowling and Ingersoll (1988)	Manual	VG1 Oval BC	1.39	429	3	N/A
Sada et al. (1996)	Manual	VG1_GRS_79	10	400	2	10
Vasavada et al. (1998)	Manual	GLL_GRS_96	10.2	Unknown	Unknown	N/A
Simon-Miller et al. (2002)	Manual	GLL_GRS_96	1.8	168	0.6	N/A
Simon-Miller et al. (2006)	Manual	GLL Oval BC	0.63	152	4	N/A
Simon-Miller et al. (2006)	Manual	CAS_OBA_00	0.88	180	2	N/A
Simon-Miller et al. (2006)	Manual	HST_OBA_06	0.68	236	4	N/A
Cheng et al. (2008)	Manual	New Horizons Oval BA	0.5	508	2	N/A
Rossow et al. (1990)	2-Pass	Pioneer Venus	4	47,000	100	N/A
Toigo et al. (1994)	1-Pass	GLL Venus	1	2588	40	N/A
Tokumaru et al. (1995)	1-Pass	VG1 West of GRS	10	3840	1	N/A
Vasavada et al. (1998)	1-Pass	GLL_GRS_96	1.2	100,000	300	N/A
Gierasch et al. (2000)	1-Pass	GLL Jovian Storm	1.55	495	2	N/A
Read et al. (2005)	2-Pass	GLL_GRS_96	1.2	1536	4	N/A
Salyk et al. (2006)	1-Pass	CAS Jupiter	9.5	200,000	3	N/A
Choi et al. (2007)	2-Pass	GLL_GRS_00	1.8	30,000	81	N/A
Del Genio et al. (2007)	1-Pass	CAS Saturn	10.5	11,398	1	2–10

using nearby velocity vectors. However, if the vector field is sparse, the closest vectors are far away and the needed approximations are not good. Also, because the arc-length actually traveled by a cloud feature can be significantly greater than the displacement in an image pair, manually-derived velocity magnitudes are systematically under-estimated.

A problem common to most of the published velocity fields that were extracted from image pairs (using manual or automated methods) is the lack of a good measure for their uncertainties. In most cases, the reported uncertainty is defined to be the pixel size divided by the time interval over which the displacement occurred. Generally, this measure is much smaller than the true error (see Section 4.1). This measure of uncertainty, based on pixel size, is misleading because it has no dependence on the quality of the image (e.g., whether the feature was at the planet's limb or whether the feature was obscured by haze), or whether there are inconsistencies in the vector field, such as cloud feature paths that cross each other. In fact, the uncertainty based on pixel size would be reported as small even in cases where the camera failed to take a picture. A good measure of uncertainty should account for all of the errors (see Section 7).

One advantage of automated methods, such as *Correlation Image Velocimetry* (CIV) (Fincham and Spedding, 1997), for extracting velocities is the large number (often, several hundred thousand) of velocity vectors they produce for features such as the GRS or Oval BA. Another advantage of automated methods is their ability to obtain sub-pixel accuracies.

However, a drawback of the published automated methods is that, in most cases, they are limited by low tolerances for noise and for poor spatial resolution. Moreover, automated methods usually fail to find any correlations when the separation between Jovian images is greater than  $\sim 2\text{ h}$  (Table 1). This is unfortunate because high precision velocity measurements usually require longer separation times (e.g., 10 h) because the uncertainties in the velocities tend to be inversely proportional to the separation time between images. This drawback becomes painfully apparent when using HST to observe Jupiter. Because the usable observation time during one HST orbit is  $\sim 45\text{ min}$ , and because a Jovian day is approximately 10 h, the image pairs of a Jovian vortex have separation times of either  $\sim 45\text{ min}$  (which produces very large uncertainties – see Section 7) or  $\sim 10\text{ h}$ . However a 10-h separation with HST images is too big for the published automated

velocity-extraction methods to find correlations (tie-points). Thus, due to our need to use HST images to track the time variability of the GRS and the Oval BA with high accuracy, we developed the ACCIV velocity retrieval method. ACCIV not only allowed us to use HST images separated by 10 h, but also allowed us to use Galileo and Cassini image pairs with large separation times. (ACCIV could also be used with navigated and deprojected image pairs from Voyager and New Horizons.)

In Section 2 we review the literature in which velocities of planetary vortices are extracted. In Section 3 we give an overview of how ACCIV works and two new methods for computing uncertainties. In Section 4 we demonstrate ACCIV, along with a manual extraction method, on synthetic data, compare their performances and compare the actual errors in the extracted velocities with our computed uncertainties. In Sections 5 and 6 we present our results for the time variability of the GRS and Oval BA, respectively. We discuss the methodology of error analysis in Section 7 and present conclusions in Section 8. Appendix A gives a detailed description of the ACCIV method, and Appendix B summarizes the sets of observations that we used with ACCIV to extract velocity fields of the GRS and Oval BA.

## 2. Review of velocity-extraction methods

Methods for extracting velocities from cloud image pairs can be divided into two categories: *automated* methods, in which a computer algorithm finds matching features by correlating small boxes of pixels in two cloud images, and *manual* methods, in which a human operator finds “tie-points” that match the location of a feature in one image to the same feature’s location in a second image. Previous extraction methods of the velocities of planetary features are summarized in Table 1 and discussed below. Table 2 summarizes the image data sets that have been used by us and others to extract velocities of Jovian vortices. The data sets that we used ourselves to extract velocity fields are described in further detail in Appendix B.

To date, the majority of Jovian velocity measurements were made manually. Manual methods were pioneered by Mitchell et al. (1981) and have been used extensively since then (Table 1). Manual methods have four advantages over automated methods. First, manual methods for deriving velocities are simple to understand and implement. Second, unlike almost all automated methods, manual methods can be used to track Jovian cloud features over time separations as large as 10 h. The human eye can usually follow the cloud motion over a 10-h period, whereas automated methods can be “confused” by cloud feature distortion. Third, a person choosing tie-points by eye can be reasonably sure that a feature has been tracked correctly (though we found that manual misidentifications do occur). Finally, automated methods produce

velocity fields that are averaged over a length scale related to the size of the correlation box – (see Section 4.2), but manual velocity fields are not averaged (unless a final correlation step is used to refine the measurements).

Automated methods for extracting velocities were originally designed for laboratory flows in which the fluid is seeded with particles that act as passive tracers. These methods are robust enough to account for the advection of particles in and out of the imaging plane (in a manner that mimics the creation and dissipation of Jovian clouds). In this work, we focus on the CIV automated method (Fincham and Spedding, 1997) because it is used as a subroutine of our ACCIV method. CIV is widely used for laboratory flows. Although the original CIV algorithm used a single pass, Fincham and Delerce (2000) added a second pass that substantially decreases systematic errors, making it possible to track features to a precision *much smaller than a pixel*. CIV with two passes uses the velocity obtained from the first pass to guide the second; a correlation box of pixels in the second image is deformed according to a linear transformation intended to take into account small deformations by the flow. We shall use the term CIV to refer to the 2-pass, rather than the 1-pass, version of the algorithm throughout the rest of this work.

Rossow et al. (1990) applied a novel, 2-pass automatic correlation method to Pioneer images of Venus. The images pairs were separated by ~4 h, and ~50,000 velocity vectors, covering most of the planet, were found. Both passes of this method are similar to the first CIV pass (i.e., the method does not attempt to distort the correlation box in either pass).

Toigo et al. (1994) developed a 1-pass method for deriving velocities from Galileo images of Venus. Their method was designed to target only the brightest and darkest cloud features. This method included several techniques for removing spurious data points and a technique for fitting the data to a smooth function composed of spherical harmonics. They found 2588 velocity vectors from 15 image pairs.

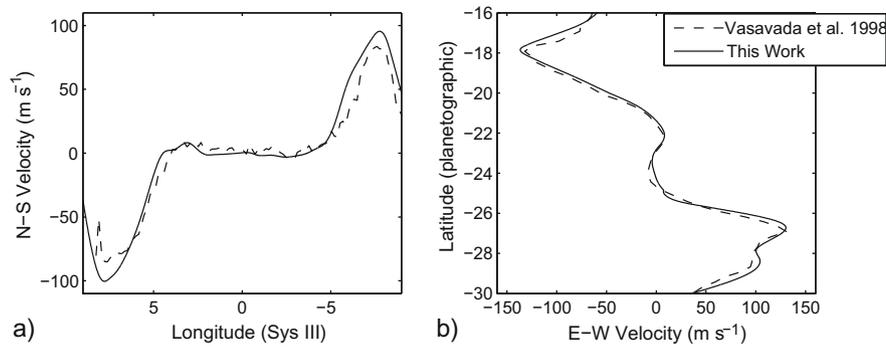
Tokumaru et al. (1995) designed a correlation method using quartic spline patches (allowing for both displacement and *non-linear* deformation of the correlation box). They found a smooth velocity field with 4000 independent degrees of freedom using Voyager 1 images of a mostly zonal region to the west of the GRS. A drawback to this method may be the very high computational cost of the minimization process required to find the best-fit parameters for each quartic spline patch.

Vasavada et al. (1998) combined a 1-pass technique, similar to the first pass of CIV, with a manual method to find a velocity field of the GRS from Galileo mosaics from June 1996 (data set *GLL\_GRS\_96*) separated by about 1 h. The quality of these mosaics is exceptional, and Fig. 5 shows that their results from 1-h separated images are similar to our ACCIV results found using 10-h

**Table 2**

A summary of the image data sets discussed in this paper. Longitudes are with respect to System III; latitudes are planetographic. The effective resolution will be defined in Section 3.4. For the synthetic data sets, the pixel and effective resolution lengths are in units where the image size is equal to unity. We did not have access to the navigated mosaics in the *VG1\_GRS\_79* data set, so it was not possible to compute their effective resolutions.

Data set ID	Feature	Spacecraft	Date	Domain bounds (long., lat.)	Pixel size (km)	Effective resolution (km)
VG1_GRS_79	GRS	Voyager 1	March 1979	(88, 68), (–30, –15)	24	Unknown
GLL_GRS_96	GRS	Galileo	June 1996	(328, 309), (–30, –15)	18	52
GLL_GRS_00	GRS	Galileo	May 2000	(12, –18), (–30, –17)	11	71
CAS_GRS_00	GRS	Cassini	December 2000	(66, 34), (–35, –11)	115	270
HST_GRS_06	GRS	HST	April 2006	(265, 233), (–35, –11)	57	163
CAS_OBA_00	Oval BA	Cassini	December 2000	(249, 227), (–46, –24)	110	270
HST_OBA_06	Oval BA	HST	April 2006	(297, 275), (–46, –24)	55	163
<i>Synth1</i>	Gaussian vortex	N/A	N/A	N/A	0.0039	0.014
<i>Synth2</i>	Gaussian vortex	N/A	N/A	N/A	0.002	0.014
<i>Synth3</i>	Gaussian vortex	N/A	N/A	N/A	0.0039	0.028
<i>SynthCusp</i>	Cusped vortex	N/A	N/A	N/A	0.002	0.0056



**Fig. 5.** Comparison of our GRS velocity field produced with ACCIV from 10-h separated images with the velocity field from Vasavada et al. (1998), produced from 1-h separated images. Both data sets were produced using images from the *GLL\_GRS\_96* data set. (a) The north–south velocity along the major diameter, (b) the east–west velocity along the minor diameter. The velocity fields are similar in most locations, but those with the 1-h separations have velocity peaks that are more rounded.

separated images from the same data set. The main difference between the results using ACCIV and the method of Vasavada et al. (1998) is that the latter rounds the velocity peaks more than ACCIV. Vasavada et al. (1998) argued that for these Galileo mosaics, manual methods were better suited for resolving fine structures in the velocity field than their automated method.

Gierasch et al. (2000) applied a version of Vasavada’s technique with multiple Galileo image pairs to study moist convection in a Jovian storm. This version used cross-correlations augmented with manual feature tracking in selected locations to produce a velocity field from images separated by about 45 min.

Read et al. (2005) were the first to show qualitatively that CIV could be used to produce a velocity field of the GRS from Galileo mosaics (data set *GLL\_GRS\_96*) with a 1.2-h separation between images, although they did not provide a quantitative analysis of the results.

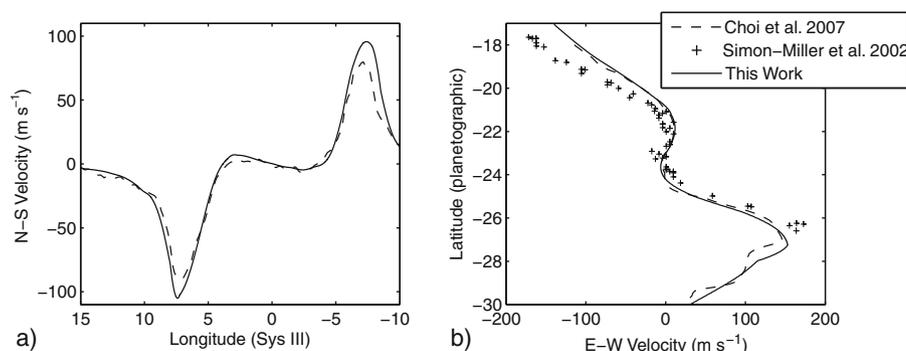
Salyk et al. (2006) applied an updated version of the 1-pass automated method of Vasavada et al. (1998) to Cassini images of Jupiter from December 2000 that show large portions of the planet. They added a technique for eliminating spurious correlations by correlating features both forward and backward in time. They were able to find correlations between images separated by 10 h. The extracted velocity field was so spatially coarse-grained that it primarily captured the zonal flows (with straight streamlines); the method missed most of the vortices and other features that would have significantly rotated and sheared the clouds over 10 h and thereby prevented the method from finding correlations.

Del Genio et al. (2007) used a technique similar to Salyk’s to extract velocities from Cassini images covering large portions of the southern hemisphere of Saturn. The method produced  $\sim 11,000$

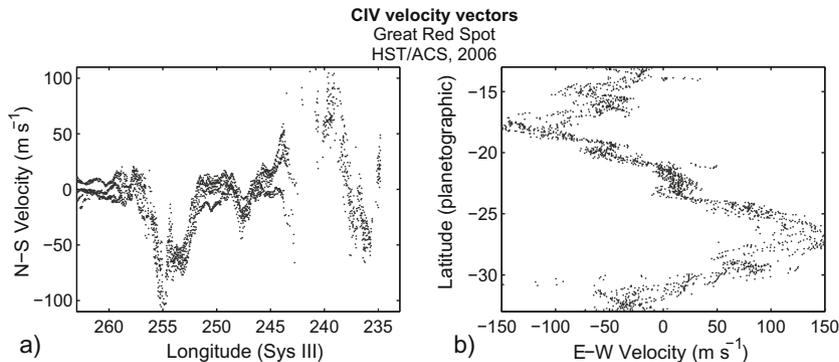
velocity vectors from images separated by  $\sim 10.5$  h. The method included a measure of uncertainty based on varying the minimum correlation value below which vectors were removed from the velocity field. Based on this measure, the authors found uncertainties of  $\sim 2\text{--}10$  m s<sup>-1</sup>. As with the method of Salyk et al. (2006), the correlation technique was successful over 10.5 h because the velocity field was dominated by zonal flows, with straight streamlines, rather than vortices.

Choi et al. (2007) designed a 2-pass method, with a 100-pixel correlation box in the first pass and a 10-pixel correlation box for the second pass. Their method did not include automated techniques to remove spurious vectors; instead they manually removed vectors that did not correlate with their neighbors. From May 2000 Galileo mosaics with 1- and 2-h separations (data set *GLL\_GRS\_00*), they produced a GRS velocity field made up of  $\sim 30,000$  vectors. Fig. 6 shows that their velocities were similar to those produced by our ACCIV method from the same mosaics, with the exception that the magnitudes of the velocity peaks from Choi et al. (2007) are smaller than those from ACCIV. The manually-derived velocities of Simon-Miller et al. (2002), using the same Galileo mosaics, have a higher velocity magnitude at the peak of the eastward-going jet stream than the fields extracted with either of the automated methods, but Fig. 6 shows the manually-derived velocities differ significantly at almost all locations from our ACCIV velocities and from the velocities extracted by Choi et al. (2007).

Note that other than the ACCIV results reported here, all of the automated methods that were successfully applied to Jovian vortices (in contrast to primarily zonal flows) used Galileo mosaics, and all used separation times less than 2 h. We tested CIV on June 1996



**Fig. 6.** As in Fig. 5, but with velocities using the mosaics taken by Galileo in May 2000 (*GLL\_GRS\_00*). The comparison is between velocities extracted with the automated method of Choi et al. (2007) using mosaics with 1- and 2-h separations and velocities extracted with ACCIV using mosaics with 10-h separations. The velocities manually-derived by Simon-Miller et al. (2002) with the same mosaics differ significantly from both of the automatically-derived velocity fields. Manual measurements were available only for the minor axis of the GRS.



**Fig. 7.** Principal axis velocities produced by CIV from the *HST\_GRS\_06* data set with a 40-min separation between images. (a) The north–south components of all velocity vectors within  $0.2^\circ$  of the east–west axis ( $23^\circ\text{S}$  planetographic latitude) of the GRS. (b) The east–west components of all velocity vectors within  $0.2^\circ$  of the north–south axis ( $248^\circ$  longitude) of the GRS. The correlation uncertainty is  $\sim 32\text{ m s}^{-1}$ .

Galileo mosaics of the GRS separated by 9–10 h and found no useful correlations within the vortex. Choi et al. (2007) found that their automated method produced a “high percentage of spurious results” when applied to the same mosaics. It seemed that the rotation and shearing of the cloud features by the vortex flow were too much for either algorithm to handle. Although the CIV results for the 1-h separated Galileo mosaics looked promising, we found that when we applied CIV to HST images of the GRS separated by about 40 min we obtained inaccurate results, as shown Fig. 7, in part because of the small time separation and in part because haze reduced cloud feature contrast close to Jupiter’s limb.

### 3. Principles of ACCIV

The ACCIV algorithm (described in detail in Appendix A) has three parts: (1) “getting started”, i.e., finding the first set of correlations (albeit, with large uncertainties) between images separated by large times, (2) iterative improvement of the velocity, and (3) determination of the uncertainties. Because the method we use to determine uncertainties can help to motivate the other two parts of ACCIV, we consider the third part first.

#### 3.1. Uncertainties

Consider the uncertainty of a temporally steady velocity field  $\mathbf{v}(x, y)$ , extracted from a pair of images taken at times  $t_1$  and  $t_2$ . Here, the velocity could have been derived by any method, including a manual method. Use any numerically accurate method to advect the pixels of the first image at time  $t_1$  forward to time  $t_2$ . The numerical advection is done by integrating the pixel intensity  $I(x, y, t)$  forward from time  $t_1$  to time  $t_2$  using the advection equation

$$\frac{\partial I}{\partial t} = -(\mathbf{v} \cdot \nabla)I. \quad (1)$$

With sufficiently small step size, the errors due to numerical advection can be made insignificant with respect to the velocity uncertainties. If the velocity field  $\mathbf{v}$  were exact, then the numerically-advected image would correlate perfectly with the pixels in the image at time  $t_2$ . If the velocity has errors, then the features of the advected image at time  $t_2$  are offset from those of the true image at time  $t_2$ . We can use any automated velocity retrieval algorithm to find – not velocities – but the offsets between the advected and true images at time  $t_2$ . For each correlation, we define this offset as the local *correlation location uncertainty*. (In practice, the offset can be measured at any time between  $t_1$  and  $t_2$ ; we typically measure the local *correlation location uncertainty* midway between these two times.)

If the separation time between two images is not too small, the *correlation location uncertainty* is a good measure of the uncertainty in the location of a dynamical feature, such as velocity peak (see Section 4). However, in the limit that the time separation between images goes to zero, the *correlation location uncertainty* goes to zero, and for separation times less than 2 h, the systematic errors are greater than the *correlation location uncertainty*, so the latter uncertainty is not a useful measure. We define the local *correlation velocity uncertainty* to be the *correlation location uncertainty* divided by the separation time between the images. (The *correlation velocity uncertainty* does not go to zero as the separation time between images goes to zero.) The *correlation velocity uncertainty* is a good measure for all image separation times of the uncertainty in the velocity due to all of the sources of random errors in the location and the magnitude of a velocity vector, including errors due to navigation, image resolution, changing lighting conditions, clouds at different elevations, and misidentified correlations (Section 4).

#### 3.2. Obtaining correlations over long time separations

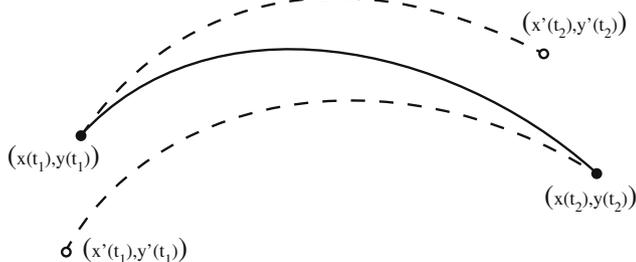
The main advantage of ACCIV is that it works with long time separations between images, while other automated methods fail to find *any* correlations (tie-points). For example, in a pair of HST images that we obtained of the GRS in April 2006 (data set *HST\_GRS\_06*) that were separated by 9.75 h, CIV found no tie-points. Within the same data set, we also had images at times  $\hat{t}_1 \equiv t_1 + 40\text{ min}$  and at  $\hat{t}_2 \equiv t_2 - 40\text{ min}$ . CIV was able to find many correlations between the two pairs of images that were separated by only 40 min. Using CIV with the pair of images at  $t_1$  and  $\hat{t}_1$ , we were able to produce a velocity field (Fig. 7) of the GRS at time  $t_1$  with large uncertainties ( $20\text{--}30\text{ m s}^{-1}$ ). Similarly, CIV produced a crude velocity field at time  $t_2$  using the image pair at  $t_2$  and  $\hat{t}_2$ .

Using the crude velocity at  $t_1$ , we advected the first image forward for  $\sim 5\text{ h}$  from  $t_1$  to the “half-way” time  $t_{12} \equiv (t_1 + t_2)/2$ . We then advected the second image backward in time from  $t_2$  for  $\sim 5\text{ h}$  to the “half-way” time  $t_{12}$  using the other crude velocity field. CIV was able to find several hundred thousand independent correlations between these two advected images at time  $t_{12}$ . These correlations at the “half-way” time then allowed us to indirectly identify the analogous correlations between the first image at time  $t_1$  and the second image at  $t_2$  by “undoing” the advection steps. Thus, although CIV could not find the correlations directly between the images at  $t_1$  and  $t_2$ , our new advection step allowed us to find the correlations indirectly.

The reason that ACCIV can find correlations, while other automated methods cannot, is the following. In the first pass of a 2-pass automated method for extracting velocities, a correlation or tie-point pair is found by starting with a “box of pixels” in the first image, and then finding a second “box of pixels”, well-correlated with the first box. In 2-pass automated methods, a first-approximation of the velocity field is found from the first pass and from that velocity field a shear field is constructed. In the second pass, sophisticated 2-pass automated methods such as CIV look for a correlation between a “box of pixels” at one location in the first images and a “box of pixels” *distorted by the shear* in the second image. Because the 2-pass method accounts for some distortion, not just a displacement, of the Jovian clouds by the velocity, the method produces more accurate results than a 1-pass method can. However, for long time separations, the pixels are not just displaced and sheared, they are distorted by the rotation and twisting due to the complex velocity field. ACCIV uses the *full*, but crude, velocity field from the previous iteration to both displace and distort the pixels *before* applying the two passes in the CIV method.

### 3.3. Iterative improvement of the velocity

The iteration scheme in ACCIV can be explained by examining the tie-points in Fig. 8. Let  $(x(t_1), y(t_1))$  be the location of a feature at time  $t_1$  obtained from the first image, and let  $(x(t_2), y(t_2))$  be the location of this same feature at time  $t_2$ , as seen in the second image. These two locations are indicated in Fig. 8 with solid black dots. Now produce a trajectory, or path of the feature, by using the extracted velocity *field* (which is made from a synthesis of the information from all the tie-points) to numerically advect the point  $(x(t_1), y(t_1))$  forward in time to its location at time  $t_2$ . This new location is indicated by the open circle labeled  $(x'(t_2), y'(t_2))$ , and its computed trajectory is indicated by the upper (dashed) path in Fig. 8. Because the velocity field used in the advection is not exact, the open circle at  $(x'(t_2), y'(t_2))$  is not coincident with the actual feature location  $(x(t_2), y(t_2))$ . (The displacement between these two points is the *correlation location uncertainty* discussed in Section 3.1.) In a similar manner, we can use the extracted velocity field to compute a trajectory backward in time, from time  $t_2$  to  $t_1$ , starting at the location  $(x(t_2), y(t_2))$  and ending at the point  $(x'(t_1), y'(t_1))$ , shown as an open circle. This latter trajectory is the lower (dashed) path in Fig. 8. Neither the upper nor the lower paths (dashed curves) in Fig. 8 is a satisfac-



**Fig. 8.** ACCIV refinement of the trajectory followed by a feature between time  $t_1$  and time  $t_2$ . The upper dashed path is the trajectory that leads from the actual (solid circle) location of the tie-point in the first image  $(x(t_1), y(t_1))$ , to its erroneous location (open circle) in the second image  $(x'(t_2), y'(t_2))$  as computed numerically by forward integration from  $t_1$  to  $t_2$  with the assumed velocity field. The lower dashed path is the trajectory that leads back from the actual location (solid circle) of the tie-point in the second image,  $(x(t_2), y(t_2))$  to its erroneous location (open circle) in the first image  $(x'(t_1), y'(t_1))$  as computed numerically by backward integration from  $t_2$  to  $t_1$  with the assumed velocity field. The distances between the actual and erroneous locations are measures of the *correlation uncertainty*. The solid path connecting the actual locations of the tie-points at times  $t_1$  and  $t_2$  is the linear interpolation of the two dashed trajectories (see Appendix A.6).

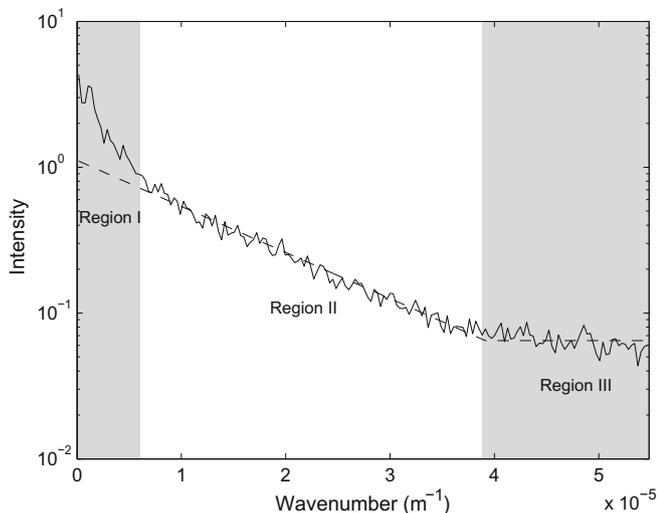
tory trajectory because neither connects the two true locations of the feature at  $(x(t_1), y(t_1))$  and  $(x(t_2), y(t_2))$ . However, it is easy to interpolate a path (solid curve in Fig. 8) that smoothly connects  $(x(t_1), y(t_1))$  to  $(x(t_2), y(t_2))$  (see Appendix A.6). Along this interpolated path, the ACCIV algorithm creates a set of new, artificial velocity vectors (equal to the numerical derivatives of the trajectory at different times between  $t_1$  and  $t_2$ ). These new vectors replace the old set of velocity vectors, and a new velocity field is created by interpolating this augmented set of velocity vectors to a uniform grid.

In order to advect numerically the tie-points and pixels it is necessary to be able to evaluate the velocity field at any arbitrary location. To do this, at the end of each iteration the velocity vectors are interpolated onto a uniform grid (approximately equal to the pixel size), and in the next iteration the velocity is interpolated from the grid to the needed spatial locations. Any interpolation algorithm that puts the velocity onto a grid must necessarily model (and therefore possibly smooth) the velocities. The *smoothing length* is defined as the characteristic length involved in the smoothing/interpolation algorithm. The velocity on the grid is used to determine the streamlines of the flow in order to find feature trajectories and to advect the cloud images. As we will show using synthetic data in Section 4.2, the smoothing length used to find the velocity field on the grid has no noticeable influence on the rounding of features in the ACCIV-derived velocity vectors. Details of how new, artificial velocity vectors are created along each trajectory and how the velocity field on the grid is constructed are in Appendix A. The advection process is iterated, each time improving the velocity field. The iteration in ACCIV should *not* be confused with the second pass in a 2-pass method. However, ACCIV does use the 2-pass CIV algorithm as a subroutine in each of its iterations.

Because the paths in Fig. 8 have finite length, they pass near the locations assigned to vectors derived from other tie-point pairs. Thus, the velocity that is used in advecting the point from  $(x(t_1), y(t_1))$  to  $(x'(t_2), y'(t_2))$ , or from  $(x(t_2), y(t_2))$  to  $(x'(t_1), y'(t_1))$ , is not only influenced by the tie-point pair  $[(x(t_1), y(t_1)), (x(t_2), y(t_2))]$ , but also by many other tie-point pairs. ACCIV starts with an initial vector field and from it creates a final, dense vector field that has smaller errors than the initial field (according to the tests in Section 4). It does *not* do this, by trying to create a “smooth” final field. If it did, and if the actual velocity field had sharp features, it would smooth the field and create large errors. Instead, what ACCIV does to create the final velocity field is to optimize the consistency of the velocity vectors. In particular, ACCIV attempts to make the curved path for each tie-point pair consistent with the velocity associate with not only that tie-point but also all the neighboring tie-points. It can be shown that this technique can preserve sharp features, even discontinuities, in the velocity field.

### 3.4. Independent vectors

One way to decrease the smoothing error created when the velocity is interpolated onto a uniform grid is to decrease the distance between the location of an extracted vector and its closest grid point. This can be done by increasing the number of extracted vectors. In CIV, the “tail” of each correlation vector or tie-point pair is placed at the center of the correlation box in the first image, and the “head” is at the center of the correlation box found in the second image. A field of correlations is created by placing the centers of the correlation boxes in the first image on a set of uniformly-spaced grid points. If a correlation could be found for each placement of the correlation box in the first image, then the number of correlation vectors created would be  $(A/\Delta^2)$ , where  $A$  is the area of the image and  $\Delta$  is the *shift distance* between grid points.



**Fig. 9.** The Fourier spectrum of an image from the *HST\_GRS\_06* data set (solid) and a piece-wise linear fit to the spectrum (dashed). The spectrum is similar to the *modulation transfer function*, the Fourier transform of the point spread function of a telescope. We distinguish three regions: I seeing halo, II near-diffraction limited core, and III the noise floor. The Fourier wavenumber of the *effective resolution length* is defined to be the wavenumber that marks the border between regions II and III. Here, the *effective resolution length* is 163 km. The maximum wavenumber in the figure corresponds to the Nyquist frequency,  $\pi/L_{\text{pixel}} = 5.5 \times 10^{-5} \text{ m}^{-1}$ , where  $L_{\text{pixel}}$  is the pixel size.

Typically, correlations cannot always be found, so the number of correlation vectors or velocity vectors is fewer than  $(A/\Delta^2)$ . The number of *independent* pieces of information in an extracted vector field is usually fewer than the number of extracted vectors because  $\Delta$  can be arbitrarily small. We define the amount of *independent* information in each image as the area of the image divided by the area of an effective resolution element, i.e., the square of the *effective resolution length* of the image. The latter length is defined in terms of the noise in the image. Fig. 9 shows a Fourier spectrum from the *HST\_GRS\_06* data set, which is a typical image of the Jovian cloud deck. The spectrum is similar to the *modulation transfer function*, the Fourier transform of the point spread function of a telescope (Sheehy et al., 2006; de Pater et al., 2004). We distinguish three regions in Fig. 9: I seeing halo, II near-diffraction limited core, and III the noise floor. The *effective resolution* of the image is defined to be the length scale at which the transition between regions II and III occurs. We use a piece-wise linear fit to the image spectrum to find the best-fit value of the *effective resolution*. The *effective resolution* of an image taken by a spacecraft depends on its technical specifications as well as the observing conditions. In general, the effective resolution length is a few pixel lengths, and its value is given in Table 2. We define the number of *independent* velocity or correlation vectors in an image pair as the amount of independent information (i.e., number of effective resolution elements) in one of the images (appropriately reduced if the correlation finder cannot find a correlation for a resolution element). When velocity extractions are performed with multiple pairs of images, the number of independent vectors is defined as the sum of the number of independent vectors in each image pair.

In general the shift distance between the centers of the correlation boxes in the first image is smaller than the length of the correlation box, so the boxes overlap. There is some “averaging” within each correlation box (although, the averaging is more complicated than an arithmetic average over all of the velocities in the correlation box – see Section 4). Even if there were a complete averaging of the information within a correlation box, the boxes can overlap and still produce independent pieces of information

(i.e., the number of independent vectors is not equal to the area of the image divided by the area of the correlation box). To see this independence, consider the analogy of a one-dimensional rolling average of a data set (e.g., the rolling averages used for daily stock market prices). The set of rolling-averaged data differs from the original set, and each point in the rolling average depends on many of the same values from the original data set as do its neighbors. Even so, the original data can be *uniquely* reconstructed from the rolling averages by applying the inverse of the “rolling average operator” (i.e., the “rolling average operation” is a linear process represented by a non-singular, and therefore invertible, matrix). In this sense, each value in the rolling average still contains unique information. Similarly, even though correlation boxes overlap, each resolution element’s extracted correlation vector contains unique information. We intentionally oversample the images (by using a shift distance smaller than the effective resolution length and also by adding velocity vectors along the arc paths in Fig. 8). We oversample to avoid interpolating (and the smoothing that goes with it) the vectors over large distances to and from the fixed grid points on which we calculate the vectors. In presenting the results of velocity extractions, we report (Tables 3–5) the number of *independent* vectors that were extracted because that number determines the spatial density of independent velocity vectors over a vortex. As shown in Figs. 1–4, the spatial density of independent vectors determines how well we can find the locations of velocity peaks and other features of the vortices.

### 3.5. Limitations of ACCIV

As noted above, ACCIV, as all correlation methods, produces velocities that have some averaging due to the finite size of the correlation box and have some smoothing due to the interpolations to and from the fixed grid. As shown in Section 4, this does not mean that the uncertainty in the locations of features found by ACCIV is large. However, in circumstances where the flow features of interest are those that vary significantly over small length scales, it may be that other methods will be more suitable than ACCIV, particularly in a case where appropriate modeling can be achieved with a relatively small number of measurements (Shetty et al., 2007). As with all cloud-tracking methods, ACCIV does not work well if the flow is time-dependent and, in particular, produces incorrect results when the velocity field changes between the first and second image. For example, the turbulent wake to the northwest of the GRS changes significantly in time. Without modeling, no cloud-tracking method works if there is significant creation or destruction of clouds between the first and second images.

Manual methods may be more suited than automated methods to detect the velocity magnitude at peaks in the flow, due to their lack of averaging. However, as shown in Section 4, due to the sparsity of vectors in hand-derived velocity fields, luck is required for a manual method to detect the peak because it may not be very likely that one of the relatively few extracted velocity vectors is situated at the peak’s location. Moreover, the locations of the peak velocities are often regions where the flow’s streamlines have small radii of curvature (e.g., the high-speed collar of the GRS). In regions of small radii of curvature, manual methods make large errors in assigning locations to the velocity vectors (see Section 4.1).

The biggest practical difficulty of ACCIV is that it needs to get “started.” That is, it requires a crude velocity field that allows the methodology in Section 3.2 to find correlations between two images separated by a long time interval. Once ACCIV finds correlations, it can bootstrap itself and iteratively improve the velocity field. However, if the clouds are so badly distorted between images that no automated or manual method can find tie-points, than ACCIV cannot get started. In addition, if the atmospheric physics

**Table 3**

A summary of velocity data described in this work for comparison with Table 1. Most velocities were extracted using the ACCIV method. Manual measurements were made in two cases for comparison. ACCIV extraction was not possible for the *HST\_OBA\_06* data set (see Section 6). The number of independent vectors was estimated using the method described in Section 3.4. The total number of ACCIV-extracted vectors is larger due to oversampling and the artificial vectors that are added along feature paths; see Appendix A.6 and Table 5. The mean density of vectors shows the relative sparsity of most previous extractions (Table 1) compared to ACCIV velocity fields. The RMS error can only be computed for synthetic data, because the exact velocity fields of the observations are not known. The definitions of the *correlation* and  $1\sigma$  uncertainties are given in Sections 3.1 and 3.6, respectively. The separation times for synthetic data are given in units of the vortex turn-around times.

Method	Data set ID or description	Sep. time hours	Number of indep. vectors	Mean vec. density ( $10^{-6} \text{ km}^{-2}$ )	RMS error (%)	Correlation uncertainty ( $\text{ms}^{-1}$ or %)	$1\sigma$ uncertainty ( $\text{ms}^{-1}$ or %)
ACCIV	GLL_GRS_96	10.2	42,000	220	N/A	2	1.2
ACCIV	GLL_GRS_00	1.8	160,000	4600	N/A	6	6.5
ACCIV	CAS_GRS_00	9.5	26,000	140	N/A	7	4.5
ACCIV	HST_GRS_06	9.75	140,000	290	N/A	5	4.5
CIV	HST_GRS_06	0.68	70,000	140	N/A	32	25.5
Manual	HST_GRS_06	9.75	443	0.4	N/A	8	4.5
ACCIV	CAS_OBA_00	9.5	4200	43	N/A	6.5	3.0
CIV	HST_OBA_06	0.68	63,000	240	N/A	9.5	5.5
ACCIV	<i>Synth1</i>	0.4	39,000	N/A	1.0%	0.8%	0.6%
ACCIV	<i>Synth2</i>	0.4	29,000	N/A	0.8%	0.6%	0.4%
ACCIV	<i>Synth3</i>	0.4	4000	N/A	2.5%	2.2%	1.9%
ACCIV	<i>SynthCusp</i>	0.2	350,000	N/A	1.6%	0.6%	0.6%
Manual	<i>SynthCusp</i>	0.2	815	N/A	5.6%	Unknown	Unknown

**Table 4**

ACCIV parameters used to produce the velocity fields for the cases in which the image separation times were short (usually 2 h or less). The image size, box size, search range, correlation box shift distance, and smoothing length are in units of pixels. The box size is the size of the correlation box for the relevant CIV pass. The search range is the range of correlation box displacements used in the first pass of CIV in each dimension (the search range for the second pass is determined automatically). The correlation box shift distance is the number of pixels by which the correlation box is shifted between each measurement. The number of image pairs is the total number of pairings of the set of images (each image can belong to more than one pairing). If the correlation box shift is greater than the effective resolution length, all correlations found by CIV are considered to be independent. Otherwise, we are oversampling, and we consider only one correlation vector per resolution element to be independent (see Section 3.4). The smoothing length is the number of pixels over which the smooth fitting function has been smoothed.

Data set	Image size	Short time intervals					Number of image pairs	Number of indep. vectors	Smoothing length
		First pass			Second pass				
		Correlation box size	Search range	Correlation box shift	Correlation box size	Correlation box shift			
GLL_GRS_96	1200 × 992	40 × 40	(−30, 30)	20	21 × 21	8	1	5185	4.3
GLL_GRS_00	3000 × 1319	50 × 50	(−60, 60)	32	30 × 30	16	2	5994	4.2
CAS_GRS_00	321 × 241	20 × 20	(−10, 10)	8	15 × 15	1	3	23,973	5.7
HST_GRS_06	641 × 481	26 × 26	(−10, 10)	24	21 × 21	4	3	69,502	12.6
CAS_OBA_00	221 × 221	20 × 20	(−10, 10)	8	15 × 15	1	1	5174	3.4
HST_OBA_06	441 × 441	26 × 26	(−10, 10)	16	21 × 21	2	6	63,339	2.9
<i>Synth1</i>	257 × 257	20 × 20	(−10, 10)	8	15 × 15	4	6	15,965	1.3
<i>Synth2</i>	513 × 513	20 × 20	(−20, 20)	8	15 × 15	4	6	11,513	1.7
<i>Synth3</i>	257 × 257	20 × 20	(−10, 10)	8	15 × 15	4	6	4898	2.4
<i>SynthCusp</i>	513 × 513	20 × 20	(−20, 20)	8	15 × 15	4	6	150,352	4

is such that the coherence time of the clouds is less than the time separation between the images, then all velocity-extraction methods that assume that clouds are passive tracers will fail. For example, in April 2006, HST images show that the clouds in the Oval BA were being created and destroyed on time scales of less than 10 h (see Section 6), and no cloud-tracking method has been able to produce velocity fields using those HST image pairs separated by 10 h.

### 3.6. $1\sigma$ uncertainty

In addition to the correlation velocity uncertainty, we found it useful to define another measure of uncertainty for cases where we do not have the original images to work with (e.g., when working with Voyager 1 manual tie-points of the GRS in Fig. 1). This measurement of uncertainty is the difference between the velocity vector at a given location and the velocity value interpolated from the grid points to that same given location. Essentially, we are using the scatter in the velocity data to measure the uncertainty. This measurement of uncertainty depends on the choice of smoothing lengths used to compute the velocity field at the grid points (see Appendix A.2). Some experimentation

indicates that this dependence is not strong as long as the smoothing length is large enough to smooth the noise in the velocity field without adding additional rounding to the velocity peaks beyond what is already present in the original data. As a measure of the “average” uncertainty of the whole data set, we take the *root mean square* (RMS) value of all the local uncertainties. We call the result the  $1\sigma$  uncertainty since it measures the standard deviation of the velocity vectors from the smoothed velocity field on the grid interpolated to the same locations. The  $1\sigma$  uncertainty tends to be slightly smaller than the correlation uncertainty (see Table 3).

## 4. Application of ACCIV to synthetic images

In order to determine the actual errors, rather than uncertainty estimates, it is necessary to know the exact velocity fields. We therefore constructed tests with two artificial data sets. The first test was designed to see how ACCIV handled images that changed over time by processes other than advection. The second test was designed to demonstrate the relative strengths and limitations of ACCIV compared with a manual method applied to a velocity field with a sharp peak.

**Table 5**

ACCIV parameters used to produce the velocity fields for the cases in which the image separation times were long (about 10 h). Parameters are listed as in Table 4. The number of artificial vectors is calculated after ACCIV places velocity vectors along each feature path.

Data set	Image size	Long time intervals									
		First pass				Second pass		Number of image pairs	Number of indep. vectors	Smoothing length	Number of artificial vectors
		Correlation box size	Search range	Correlation box shift	Correlation box size	Correlation box shift					
GLL_GRS_96	1200 × 992	100 × 100	(−25, 25)	20	40 × 40	2	2	41,854	8.5	664,971	
GLL_GRS_00	3000 × 1319	100 × 100	(−25, 25)	8	40 × 40	2	3	160,575	42.3	6,586,591	
CAS_GRS_00	321 × 241	20 × 20	(−10, 10)	8	15 × 15	1	4	26,232	11.3	1,069,209	
HST_GRS_06	641 × 481	26 × 26	(−20, 20)	8	21 × 21	2	16	140,138	9.6	2,135,339	
CAS_OBA_00	221 × 221	20 × 20	(−10, 10)	8	15 × 15	1	1	4232	5.7	194,780	
Synth1	257 × 257	20 × 20	(−10, 10)	8	15 × 15	4	16	38,814	3.3	305,699	
Synth2	513 × 513	40 × 40	(−20, 20)	16	30 × 30	8	16	29,006	2.4	226,224	
Synth3	257 × 257	20 × 20	(−10, 10)	8	15 × 15	4	16	4036	2.8	102,589	
SynthCusp	513 × 513	40 × 40	(−20, 20)	16	20 × 20	2	16	345,486	4	2,708,608	

#### 4.1. First test: ACCIV applied to synthetic images of particles in a Gaussian vortex

This test not only verifies the accuracy of the ACCIV method, but also compares the *correlation velocity uncertainty* (see Section 3.1) to the actual errors in the extracted velocity field. In addition, this test enables us to explore the sensitivity of ACCIV both to the properties of the input images – pixel resolution, *effective resolution*, time between images, etc. – and to the user-specified parameters that control ACCIV (correlation box size, smoothing length, etc.). To carry out the test, we construct a series of synthetic images of particles that are passively advected by an analytic flow field. The exact velocity field used in each data set in this test (*Synth1*–*Synth3* in Table 2) has a Gaussian radial distribution of vorticity with full-width-at-half-maximum  $r_0$  and velocity  $v_0$  at  $r = r_0$  (see gray curve in Fig. 10):

$$v_\phi(r) = v_0 \left(\frac{r}{r_0}\right)^{-1} \left(\frac{1 - e^{-r^2/(2r_0^2)}}{1 - e^{-1/2}}\right). \quad (2)$$

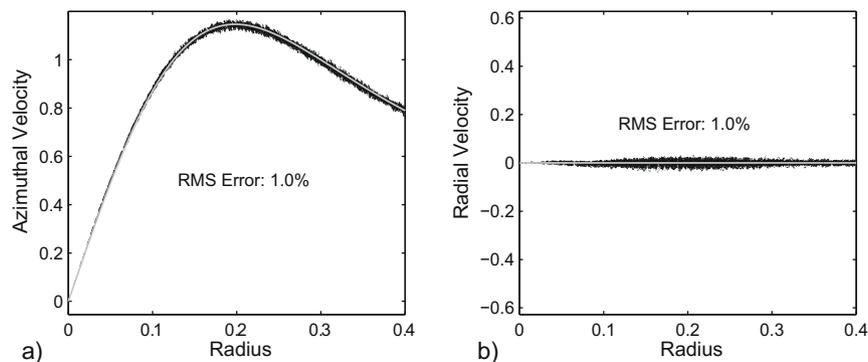
The particles were represented by radial Gaussian functions in brightness where the full-width-at-half-maximum is the particle radius. We give the images a height and width equal to 1. In these units,  $r_0 = 1/8$ . Our unit of time is chosen so that  $v_0 = 1$ . In these units, the peak velocity is  $v_{\text{peak}} \sim 1.1$ .

The top of Fig. 11 shows two images from the *Synth2* data set separated by half the turn-around time of the vortex (the time it takes a particle to travel half-way around the vortex at the radius where the velocity is a maximum), so that the curvature of particle paths is significant. Although the particles are ad-

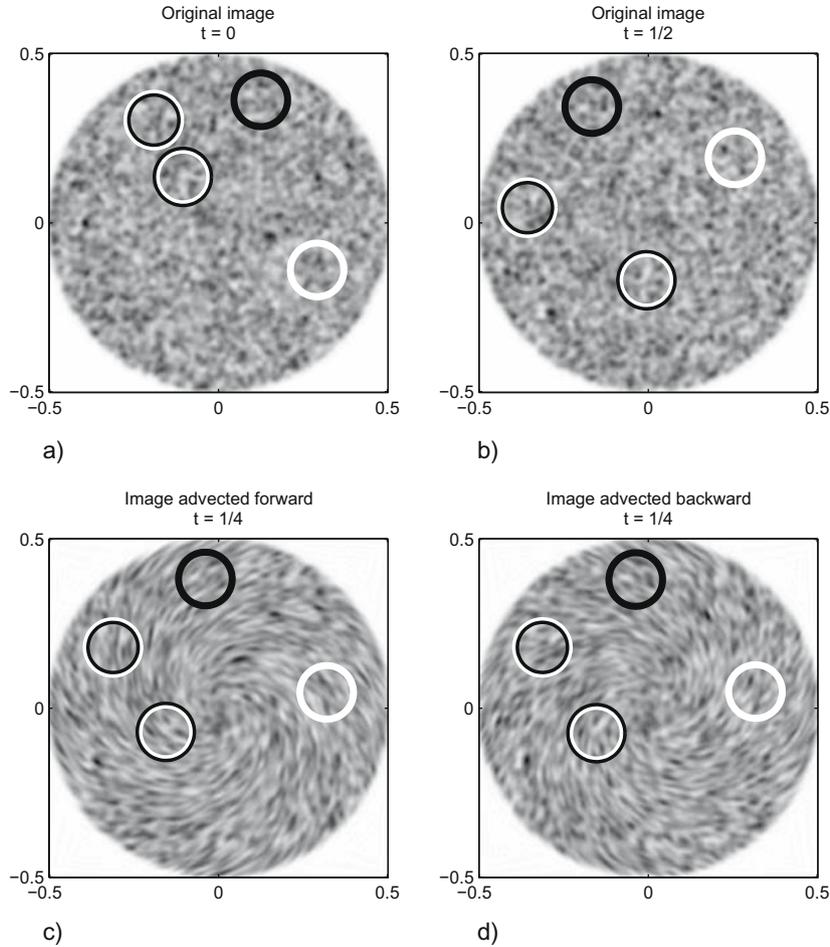
vected by the flow, the patterns formed by neighboring particles are not preserved, mimicking changing cloud patterns. In fact, the patterns change so much between the images in Fig. 11 that we were only able to find a few dozen manual tie-points. We have highlighted four features in each image to demonstrate the difficulty in correlating features in the original images, and the relative ease of finding matches in the advected images produced by ACCIV.

We used ACCIV to produce a velocity field from eight images of the synthetic flow: four taken early in the simulation at times 0, 0.012, 0.025 and 0.037 (similar to Fig. 11a), and four more taken half a turn-around time later at times 0.368, 0.380, 0.393 and 0.405 (similar to Fig. 11b). We use three different sets of synthetic images, *Synth1*, *Synth2* and *Synth3* in Table 2, to study the dependence of ACCIV on both the pixel size in the image and on the effective resolution of the image (in this case, easily changed by varying the particle size). In the *Synth1* case, the pixel size is  $1/256$  and the particle radius is  $1/128$  in our units. The radial and azimuthal components of the  $\sim 39,000$  independent velocity vectors extracted by ACCIV from the *Synth1* data set are shown in Fig. 10 along with the analytic solution. Because we know the exact solution, we can compute the *error*, defined to be the difference between the exact and extracted velocities. The RMS error ( $L_2$  norm) of the whole data set is 1.0% of the peak velocity, and the maximum error at any location ( $L_\infty$  norm) is 3.1%. The correlation velocity uncertainty of 0.8% slightly underestimates the RMS error, probably because the RMS error includes systematic errors that are not measured by the correlation velocity uncertainty (see Section 7.4).

Data set *Synth2* explores the effect of varying the images' pixel size, while keeping their effective resolution fixed. In this data set,



**Fig. 10.** (a) Azimuthal velocity field of a Gaussian vortex (*Synth1* data set). (b) Radial component of the same velocity field. In both plots, the velocities extracted with ACCIV are shown in black and the analytic solution in gray. The maximum error of the velocity field is less than 3.1% of the maximum velocity, and the RMS error is  $\sim 1.0\%$  of the maximum velocity.



**Fig. 11.** Four images from the ACCIV process applied to a synthetic data set (*Synth2*) of “particles” being advected by a Gaussian vortex. (a) and (b) Unadvected synthetic image separated by about half the turn-around time of the vortex ( $t = 0$  and  $t = 1/2$ , respectively). Finding correlations by hand between these images is difficult. The bottom panels are the images from the top panels advected forward (c) and backward (d) to the midpoint in time ( $t = 1/4$ ). It is easier to see correlations between these two images. Four features are highlighted with circles in each of the images to showcase the difficulty of finding these matches in the original images and the relative ease of finding matches in the advected images. ACCIV is able to find correlations between figures (a) and (b) without difficulty.

we reduced the pixel size to  $1/512$  while maintaining the particle size at  $1/128$ . We found that the errors decrease slightly with decreased pixel size: the RMS error is 0.8%, while the maximum error is 2.9% for the *Synth2* data set. The slight improvement probably results from the fact that the correlation box used by CIV in this run contains four times more pixels as the previous run, meaning that the maximum of the correlation function can be found more accurately. The correlation velocity uncertainty is 0.6%, again an underestimate of the RMS error.

The *Synth3* data set has the same pixel size as *Synth1*,  $1/256$ , but a particle radius of  $1/64$  giving the data set a bigger effective resolution length. Doubling the particle size (and therefore doubling the effective resolution length) increases the average error to 3.0% and the maximum error to 11.6%. In this test, the correlation velocity uncertainty is 2.2%, a fairly significant underestimate of the RMS error. The results of the three tests suggest that the error is only a weak function of the pixel resolution but a much stronger function of the effective image resolution.

In summary, the results of applying ACCIV to *Synth1* – *Synth3* data sets show that ACCIV accurately extracts velocity fields from the same image pairs from which a manual method produced only a few dozen measurements. For example, ACCIV extracted  $\sim 29,000$  independent velocity vectors from the *Synth2* images for which a manual method extracted only a few dozen. The accuracy of the extracted velocity field is limited by the images’ effective resolutions rather than their pixel sizes. Using ACCIV, we are able to pro-

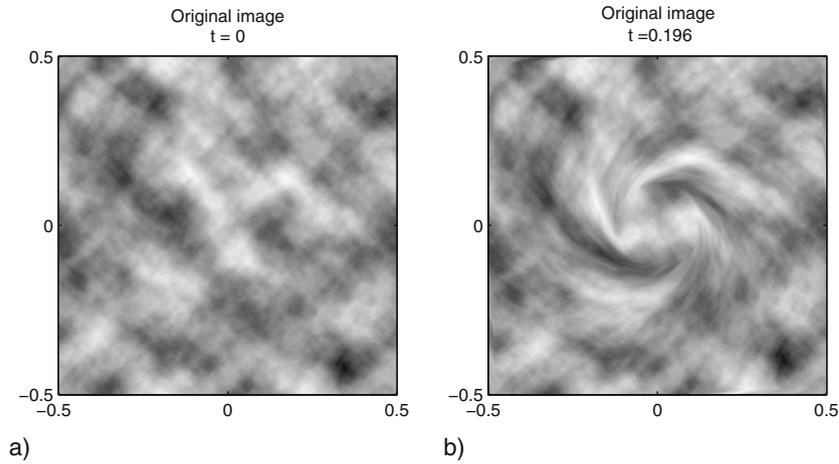
duce a velocity field on a grid with less than 1% average error, a feat that would not be possible with a manual method.

#### 4.2. Second test: ACCIV applied to synthetic images in a cusped vortex

The second test explores ACCIV’s ability to accurately find velocity fields that contain sharp cusps (similar to the high-speed collar of the GRS; see Fig. 12). In particular, we are interested in exploring the variation of the velocity error when we vary the correlation box size and the smoothing length in the ACCIV algorithm. The *SynthCusp* image data set is composed of several early frames of the simulation and several frames near the end. We use these images with short separation times to get an estimated velocity field using CIV. Then, we apply ACCIV to the synthetic images to see how well we reconstruct the original velocity field, especially near the cusp.

To generate the synthetic images, we began with a Perlin noise pattern (Perlin, 2002) that roughly resembled cloud features, and which contained information at all scales from the size of the image down to the pixel size. We advect the image forward in time using the velocity field for a constant vorticity circular vortex, defined by

$$v_{\phi}(r) = \begin{cases} v_0 \left( \frac{r}{r_0} \right) & r < r_0 \\ v_0 \left( \frac{r_0}{r} \right) & r \geq r_0 \end{cases} \quad (3)$$

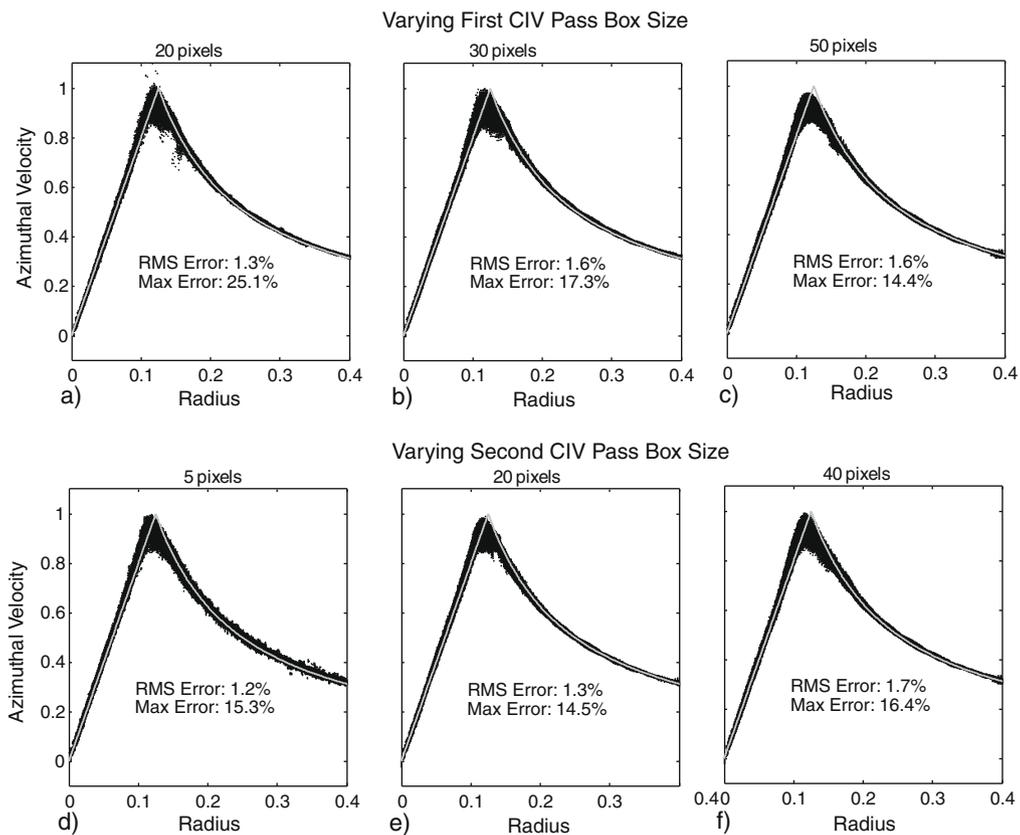


**Fig. 12.** Synthetic cloud images used to test the ACCIV method on a vortex with a sharply peaked velocity field (*SynthCusp* data set). (a) and (b) The synthetic cloud images at times  $t = 0$  and  $t = 0.196$ , respectively.

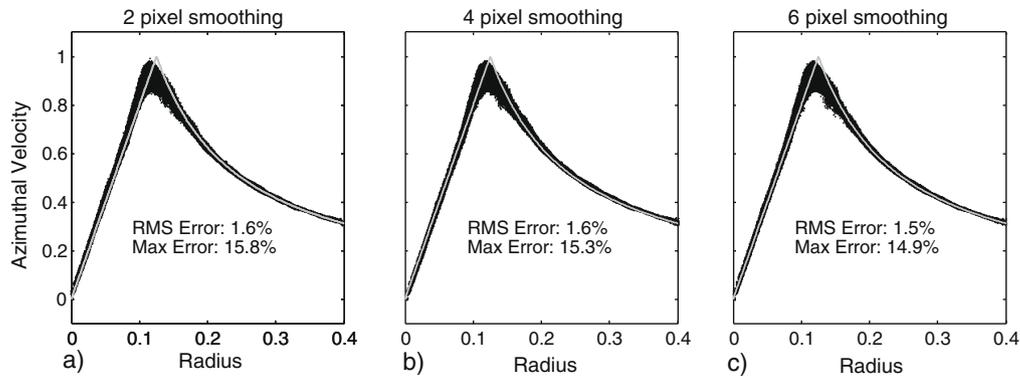
In our units,  $v_0 = 1$ , the pixel size is  $1/512$  and  $r_0 = 1/8$  (or 64 pixels). The noise pattern was advected for  $t \sim \pi/16$ , about 1/4 of the vortex turn-around time. Over this amount of time, some of the paths followed by the cloud feature paths are significantly curved, but the distortion of the original cloud image is not so great that features cannot be identified manually. We performed ACCIV velocity extraction on eight synthetic cloud images, four images at  $t = 0, 0.012, 0.025,$  and  $0.037$  (similar to Fig. 12a) and four

more taken 1/4 of a vortex turn-around time later at  $t = 0.184, 0.196, 0.209,$  and  $0.221$  (similar to Fig. 12b). To give a better idea of the ACCIV process, Appendix A.8 describes in detail its application to the *SynthCusp* data set.

The way in which the different elements of a cloud pattern within a correlation box are weighted is more complex than a simple averaging. The width of the correlation box is the theoretical upper bound of the averaging length due to the finite box size,

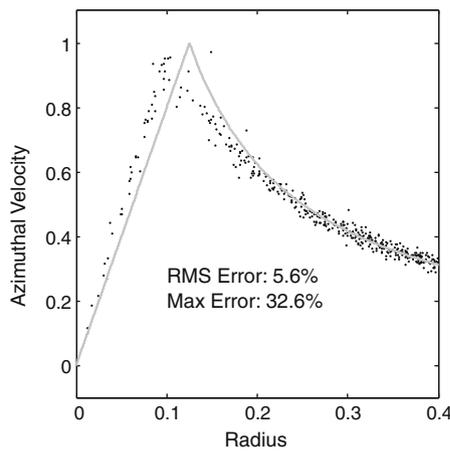


**Fig. 13.** The effects of varying the correlation box size used in the first and second passes of CIV. Each panel shows the azimuthal velocity (gray curve) of a constant vorticity circular vortex (*SynthCusp* data set) superposed with the ACCIV-derived azimuthal velocity (black points) as functions of vortex radius. In panels (a), (b) and (c), the correlation box size in the first pass of CIV was 20, 30 and 50 pixels, respectively, while the correlation box size for the second CIV pass was fixed at 30 pixels. In panels (d), (e) and (f), the correlation box size in the first pass of CIV was fixed at 40 pixels, while the correlation box size for the second CIV pass was 5, 20 and 40 pixels, respectively. The panels show that the RMS errors and the amount of rounding of the velocity peak are nearly independent of the correlation box size. The maximum error can become large when the correlation box (especially in the first pass) is too small.



**Fig. 14.** Azimuthal velocities of a constant vorticity circular vortex (*SynthCusp* data set) as in Fig. 13. The three panels illustrate the effect of the *smoothing length*, which is 2 pixels in panel (a), 4 pixels in panel (b) and 6 pixels in panel (c). The panels show that the errors (RMS or max) and the amount of rounding of the velocity peak are nearly independent of the *smoothing length*.

but results from our simulated flow show that averaging, in this case, is not strongly dependent on the length of the correlation box in either the first or the second pass of CIV. We varied the

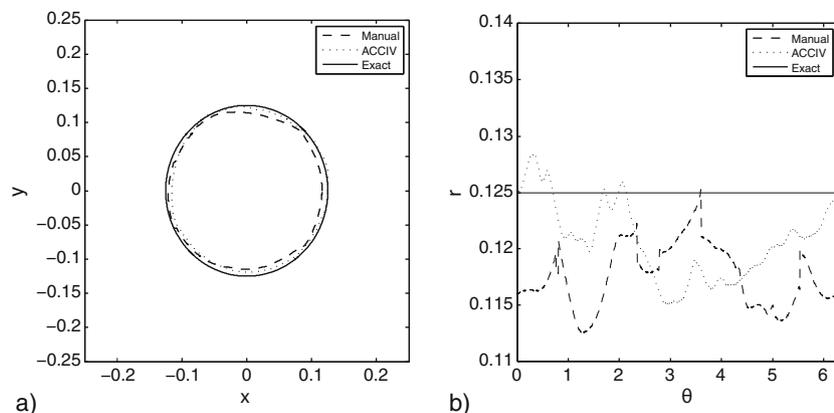


**Fig. 15.** Azimuthal velocity of a constant vorticity circular vortex together with the azimuthal component of the manual velocity field (based on the *SynthCusp* data set) as functions of vortex radius. The manual method systematically underestimates both the radius and the velocity magnitude of the velocity peak. The RMS error of the manually-extracted velocity is 5.6% and is significantly larger than the  $\sim 1.6\%$  RMS error of the ACCIV-extracted velocity.

box length in the first pass between 20 and 50 pixels, and the box length in the second pass between 5 and 40 pixels. Fig. 13 shows that the box size has little effect on the extracted velocities and their  $L_2$  error. The  $L_\infty$  error was also largely independent of the box size, except when the correlation box was so small that false matches became relatively frequent (Fig. 13a). Note that a correlation box with a length of 40 pixels is  $\sim 2/3$  the radius of the vortex, yet has surprisingly little effect on the rounding of the velocity peaks.

Our results also show that the rounding of the velocity peaks is not due to the smoothing length. We tested the effects of the smoothing length by using the same synthetic flow image as in Fig. 13. The results in Fig. 14 shows that the velocity and its  $L_2$  and  $L_\infty$  errors are nearly independent of smoothing length as it is varied from 2, 4 and 6 pixels.

To compare ACCIV with manual methods, we identified tie-points manually from two images of the synthetic flow at  $t = 0$  and  $t = 0.184$  using Matlab software (based on the *cpselect* routine). The manually-extracted velocities along with their exact values are shown in Fig. 15. The locations of the manually-produced tie-points have uncertainties greater than one pixel. This is disconcerting because many previous analyses using manual methods assumed accuracy of one pixel (Mitchell et al., 1981; Sada et al., 1996; Vasavada et al., 1998; Simon-Miller et al., 2002; Simon-Miller et al., 2006; Cheng et al., 2008). The manually-extracted velocity systematically under-predict both the radial



**Fig. 16.** Contours of maximum velocity magnitude along rays radiating out from the center of a constant vorticity vortex (the *SynthCusp* data set). The exact contour is a circle (solid line), while the ACCIV-derived contour is shown with a dotted line and the manually-derived contour is shown with a dashed line. (a) The contours in the  $x$ – $y$  plane. (b) The same contours as in (a) but shown in polar coordinates, illustrating that the ACCIV velocity field generally gives a more accurate location of velocity peaks. The mean errors in the radius of the ACCIV and manual contours were 3.6% and 5.9%, respectively.

location and magnitude of the peak velocities because streamline curvature is not taken into account. These systematic errors can be reduced if the separation time between images is shortened, but reducing the separation time increases the random uncertainties for all methods. Manual methods that account for streamline curvature exist (Simon-Miller, A., personal communication; Hueso, R. and Sánchez-Lavega, A., personal communication), but to our knowledge, neither these methods nor the results obtained with them have been published, leaving us without a means for comparing these methods with ACCIV.

Figs. 13 and 14 show that ACCIV systematically under-predicts the peak velocities but accurately finds the locations of the peaks. Fig. 16 confirms that ACCIV is better than a manual method for finding the locations of velocity peaks. The figure shows the radial location of the peak azimuthal velocity as a function of angle around our synthetic vortex using the velocities extracted with ACCIV and manual methods. The  $L_2$  and  $L_\infty$  errors in the radial location of the velocity peak are only 3.6% and 6.4%, respectively, for the ACCIV-extracted velocities, while they are 5.9% and 9.9% for the manually-extracted velocities.

### 5. ACCIV applied to the GRS

We extracted the velocity field of the GRS at four times: June 1996 from Galileo images (data set *GLL\_GRS\_96*), May 2000 from Galileo images (data set *GLL\_GRS\_00*), December 2000 from Cassini images (data set *CAS\_GRS\_00*), and April 2006 from HST images (data set *HST\_GRS\_06*). Table 2 lists the data sets, which are described in Appendix B. Figs. 17–19 and 4, respectively, show the velocities of the GRS along its two principal axes at these four times. ACCIV extracted ~42,000 independent velocity vectors from the Galileo mosaics from 1996. The velocities have extraordinarily

small correlation and  $1\sigma$  uncertainties of  $\sim 2 \text{ m s}^{-1}$  and  $\sim 1.2 \text{ m s}^{-1}$ , respectively. ACCIV extracted ~160,000 independent velocity vectors from the Galileo mosaics from 2000, with correlation and  $1\sigma$  uncertainties of  $\sim 6 \text{ m s}^{-1}$  and  $\sim 6.5 \text{ m s}^{-1}$ , respectively. The mosaics did not cover the northernmost part of the GRS, so the location of its northern boundary (and the length of its minor axis) could not be determined. ACCIV extracted ~26,000 independent velocity vectors from Cassini images from 2000, with correlation and  $1\sigma$  uncertainties of  $\sim 7 \text{ m s}^{-1}$  and  $\sim 4.5 \text{ m s}^{-1}$ , respectively. The HST image from 2006 yielded ~140,000 independent velocity vectors and had correlation and  $1\sigma$  uncertainties of  $\sim 5 \text{ m s}^{-1}$  and  $\sim 4.5 \text{ m s}^{-1}$ , respectively. These measures of uncertainty are discussed further in Section 3.1. Table 4 summarizes the input parameters and results for each CIV velocity field of the GRS using images separated by short times; Table 5 lists similar ACCIV parameters and results based on images separated by long times.

Figures 17, 18, 19 and 4 are summarized in Table 6 and Fig. 20, which show that to within the correlation velocity uncertainties (depicted with error bars), the magnitudes of the peak velocities of the GRS (with the possible exception of the eastward-going velocity peak at the southern edge of the GRS) did not change between 1996 and 2006. To within the correlation location uncertainties (also depicted as error bars), the length of the minor axis of the GRS (defined as the distance along the north–south principal axis between the locations of the east–west velocity peaks) did not change. However, the length of the major axis (defined as the distance along the east–west principal axis between the locations of the north–south velocity peaks) has shrunk monotonically between 1996 and 2006, and the decrease is greater than the correlation location uncertainties.

The contour indicating the location of the local peak velocities in the high-speed collar around the GRS is not well-approximated

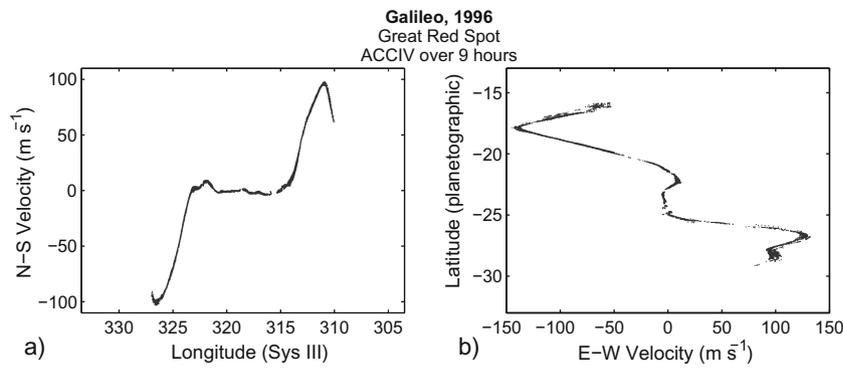


Fig. 17. Principal axis velocities of the GRS from Galileo mosaics from June 1996 (data set *GLL\_GRS\_96*). Velocity vectors were computed using ACCIV from mosaics separated by about 9 h. Only points within  $0.2^\circ$  of the principal axes were used. The correlation uncertainty in the velocity vectors is  $\sim 2 \text{ m s}^{-1}$ .

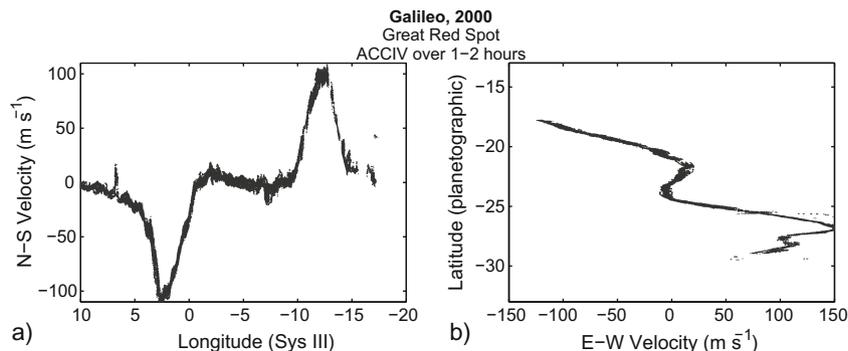
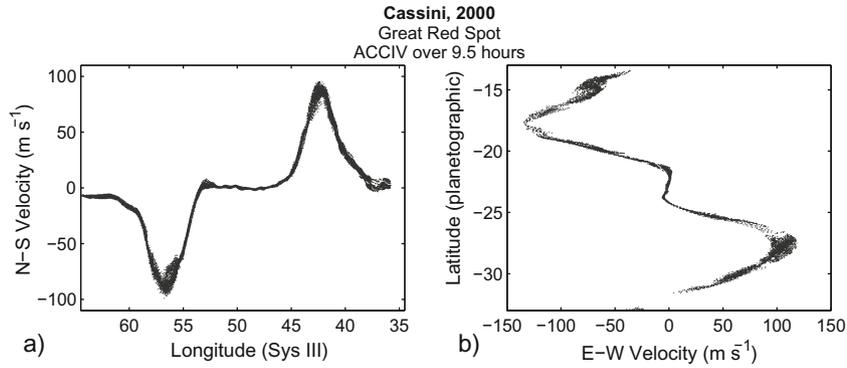


Fig. 18. Principal axis velocities of the GRS from Galileo mosaics from May 2000 (data set *GLL\_GRS\_00*). Velocity vectors were computed using ACCIV from mosaics separated by about 1–2 h. Only points within  $0.2^\circ$  of the principal axes were used. The correlation uncertainty in the velocity vectors is  $\sim 6 \text{ m s}^{-1}$ .



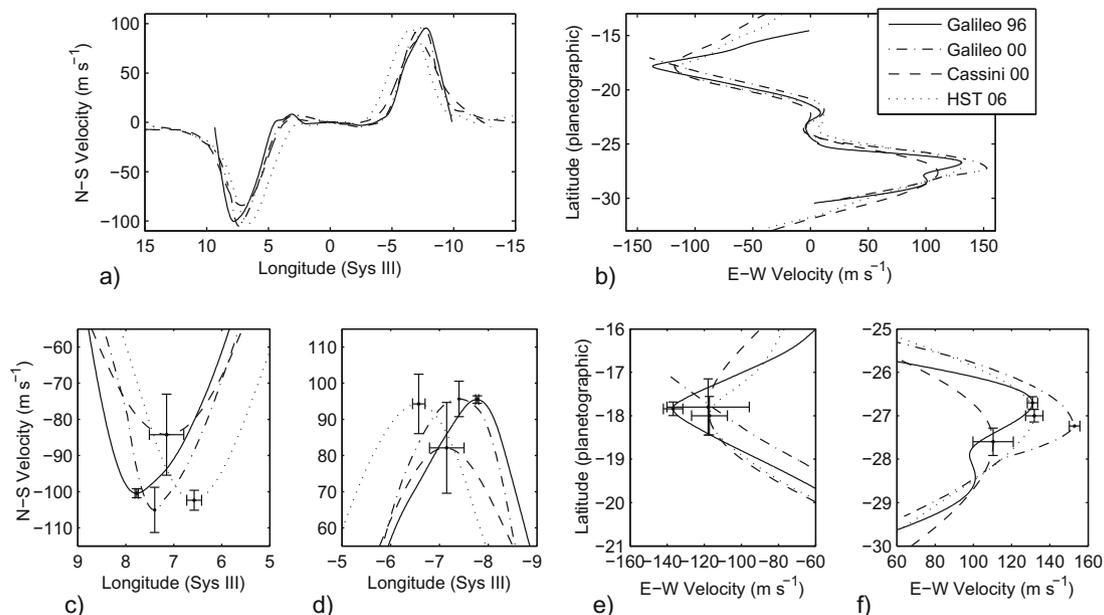
**Fig. 19.** Principal axis velocities of the GRS from Cassini images from December 2000 (data set CAS\_GRS\_00). Velocity vectors were computed using ACCIV from images separated by about 9.5 h. Only points within 0.2° of the principal axes were used. The correlation uncertainty in the velocity vectors is  $\sim 7 \text{ m s}^{-1}$ .

**Table 6**  
The major and minor diameters of the GRS over the decade spanning 1996–2006, as determined by the locations of the velocity peaks along the principal axes. Latitudes are planetographic. The northern peak in the GRS’s velocity was not captured in the Galileo mosaics from 2000, so the minor diameter at that time is unknown. These data show that the major diameter of the GRS steadily decreased from 1996 to 2006, while the length of the minor diameter fluctuated. The uncertainties listed in the Table are the RMS average of the magnitudes of all of the correlation location uncertainties of all of the independent tie-points within 500 km of each peak. The GLL\_GRS\_00 data set has image separation times less than 2 h, so as discussed in Section 3, we cannot use the correlation location uncertainty as a measure of the uncertainty of the GRS’s major diameter for this data set.

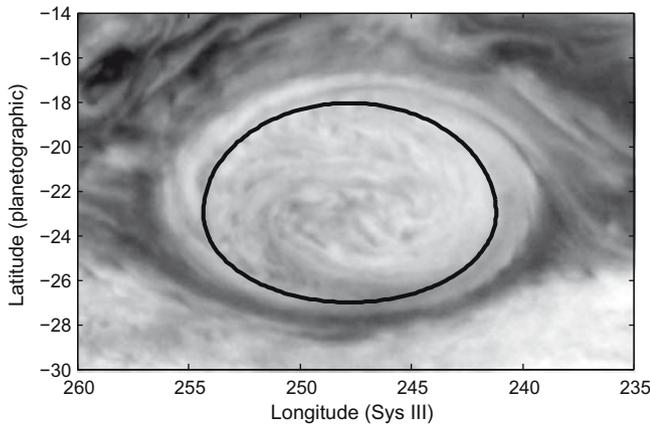
Data set	Date	Major diameter (° long.)	Major diameter (km)	Northern peak lat.	Southern peak lat.	Minor diameter (km)
GLL_GRS_96	June 1996	$15.5 \pm 0.07$	$18,100 \pm 80$	$-17.8 \pm 0.16$	$-26.7 \pm 0.15$	$9900 \pm 340$
GLL_GRS_00	May 2000	14.8	17,100	Unknown	-27.2	Unknown
CAS_GRS_00	December 2000	$14.3 \pm 0.65$	$16,600 \pm 830$	$-17.8 \pm 0.65$	$-27.6 \pm 0.31$	$11,000 \pm 1100$
HST_GRS_06	April 2006	$13.2 \pm 0.28$	$15,300 \pm 330$	$-18.1 \pm 0.44$	$-27.0 \pm 0.14$	$10,000 \pm 660$

by an ellipse because its shape is not mirror-symmetric about its east–west principal axis. However, the location of the collar is well-approximated by the union of two half-ellipses as in the black contour in Fig. 21. The upper part of the contour is a half-ellipse that passes through three points: the locations of the two velocity extrema on the east–west principal axis and the maximum of the westward-going velocity on the north–south principal axis. The lower part of the contour is a half-ellipse that also passes through

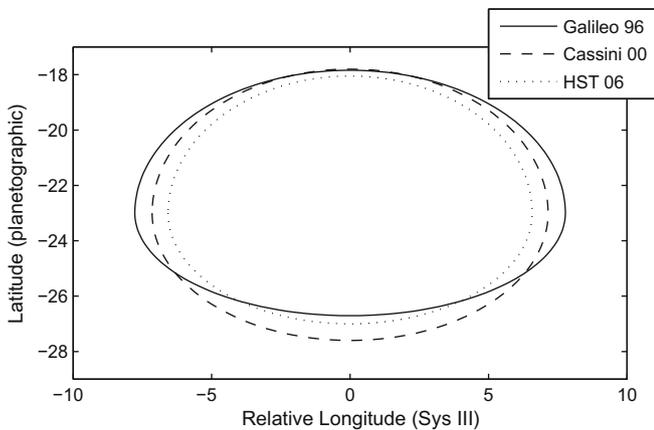
three points: the locations of the two velocity extrema on the east–west principal axis and the maximum of the eastward-going velocity on the north–south principal axis. The area circumscribed by the high-speed collar (which Shetty et al. (2007) showed to be nearly coincident with the potential vorticity anomaly of the GRS) decreased monotonically between 1996 and 2006, despite the fact that during that same interval of time, the length of the north–south principal axis increased before it decreased, if it chan-



**Fig. 20.** Curves fit through the velocities on the (a) east–west and (b) north–south principal axes of the GRS. (c), (d), (e) and (f) close-up views of the peaks with error bars indicating the uncertainties in peak velocity magnitudes and locations (see Section 5 and Table 6). The data show a steady shrinking of the GRS’s major diameter (c and d) but no trend (and perhaps no change at all) in its minor diameter (e and f). For the Galileo images taken in 2000, the correlation location uncertainty is not a useful measure of the uncertainty in the peak location due to the small separation time in the images.



**Fig. 21.** The cloud reflectivity of the GRS at 658 nm from the *HST\_GRS\_06* data set. The black contour denotes the approximate location of the GRS's high-speed collar and is the dotted contour in Fig. 22. This figure shows that the area of cloud cover is greater than the area circumscribed by the collar.



**Fig. 22.** Contours that approximate the locus of the high-speed collar of the GRS as its area decreased between June 1996 and April 2006. The upper and lower halves of each contour are half-ellipses that pass through the velocity peaks along the major and minor axes of the GRS. The major diameter of the GRS decreased during this period, while the length of the minor diameter fluctuated but remained nearly constant, showing that the GRS is both losing area and becoming rounder.

ged at all. This decrease can be seen in Fig. 22, which shows the contours of the location of the high-speed collar at the three available times: June 1996, December 2000 and April 2006. The areas of the cloud cover of the GRS (Simon-Miller et al., 2002) also de-

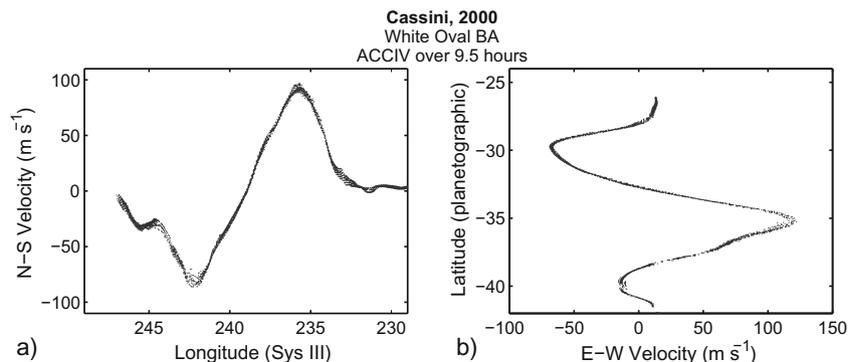
creased over this time, but Fig. 21 shows that the area covered by the clouds is not the same as the area circumscribed by the high-speed collar, and is, in fact, larger.

### 6. ACCIV applied to the Oval BA

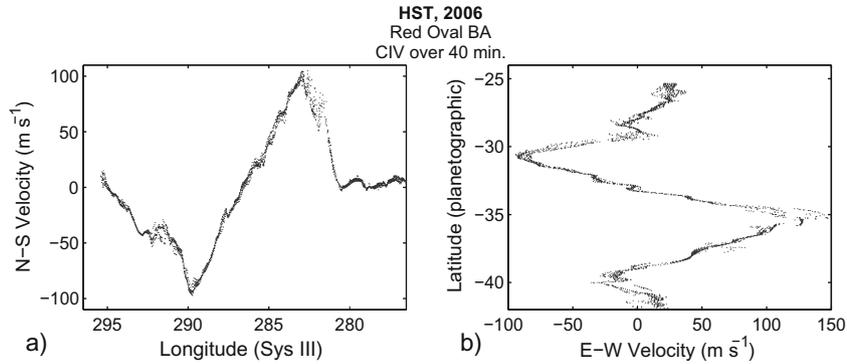
Jupiter's Oval BA was created by the merger of three White Ovals between 1998 and 2000. Therefore, the only navigated image pairs that were available to analyze were from December 2000 (Cassini data set *CAS\_OBA\_00*) and from April 2006 (HST data set *HST\_OBA\_06*). While ACCIV was able to extract velocity vectors from a Cassini 2000 image pair with a 9.5-h separation in time, ACCIV was not able to extract any tie-points from the HST 2006 image pair that had a separation time of 9.75 h. No automated method or manual method, used by us or by others, has been able to extract tie-points from this long time separation HST image pair. A visual examination of the HST images from April 2006 clearly shows that the clouds in the Oval BA (but not the GRS) had coherence times less than 9.75 h. Thus, the image pair from the 2006 HST observation (data set *HST\_OBA\_06*) used for velocity extraction by CIV had only a 40 min time separation.

Figures 23 and 24 show the velocities along the principal axes of the Oval BA during 2000 and 2006, respectively. ACCIV extracted ~4200 independent velocity vectors from the images from December 2000 with correlation and  $1\sigma$  uncertainties of  $\sim 6.5 \text{ m s}^{-1}$  and  $\sim 3.0 \text{ m s}^{-1}$ , respectively. CIV extracted ~63,000 independent velocity vectors from the images from 2006 with correlation and  $1\sigma$  velocity uncertainties of  $\sim 9.5 \text{ m s}^{-1}$  and  $\sim 5.5 \text{ m s}^{-1}$ , respectively. Table 4 summarizes the input parameters and results for each CIV velocity field of the Oval BA; Table 5 lists the ACCIV parameters and results based on the Cassini images separated by long times. The uncertainties in the CIV-extracted velocities are extraordinarily small considering that the time separation between images is only 40 min. If the accuracies of the lengths of the tie-point vectors were limited by the effective image resolution of  $\sim 163 \text{ km}$ , then the uncertainty in the extracted velocity field would be  $\sim 68 \text{ m s}^{-1}$ , comparable to the uncertainties of the velocities manually-derived from this image pair (Simon-Miller et al., 2006 see below). CIV achieves its very small velocity uncertainties because its tie-point vectors have sub-pixel accuracy. Because cloud feature paths in the Oval BA have negligible curvature over 40 min, application of ACCIV over this short separation time did not reduce the error.

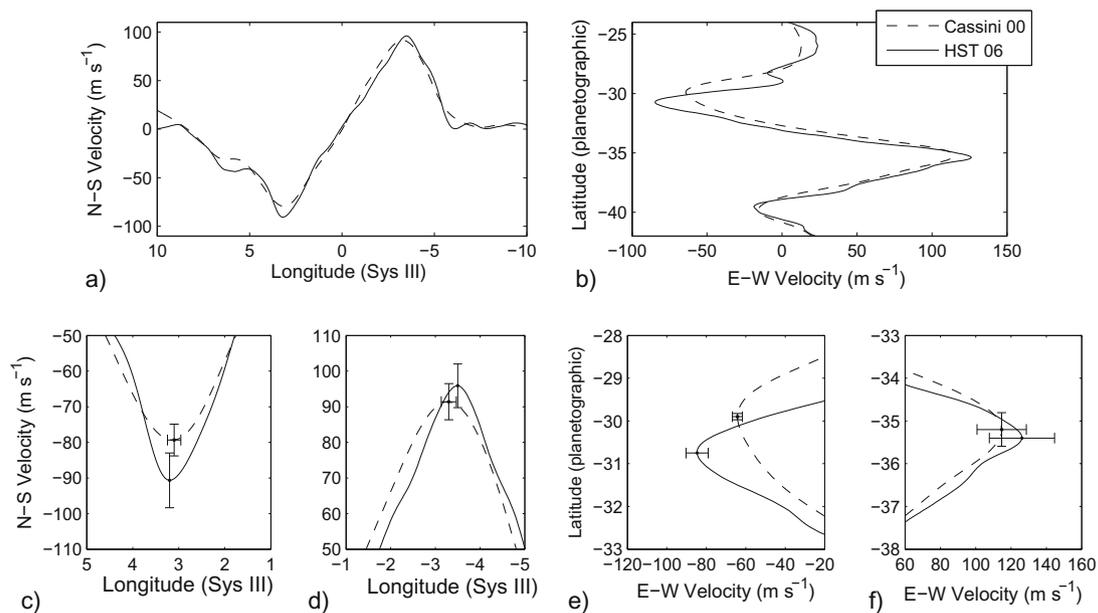
Our key results for the Oval BA are that neither its velocities nor its area changed from 2000 when it was *white* to 2006 when it was *red*. Fig. 25 shows smooth curve fits to the north–south velocities from 2000 and 2006 along the east–west axis. The two curves are identical, within the uncertainties, at almost all locations, and a similar comparison of the velocities along any axis passing



**Fig. 23.** Principal axis velocities of the White Oval BA from Cassini images from December 2000 (data set *CAS\_OBA\_00*). Velocity vectors were computed using ACCIV from images separated by about 9.5 h. Only points within  $0.2^\circ$  of the principal axes were used. The correlation uncertainty in the velocity vectors is  $\sim 6.5 \text{ m s}^{-1}$ .



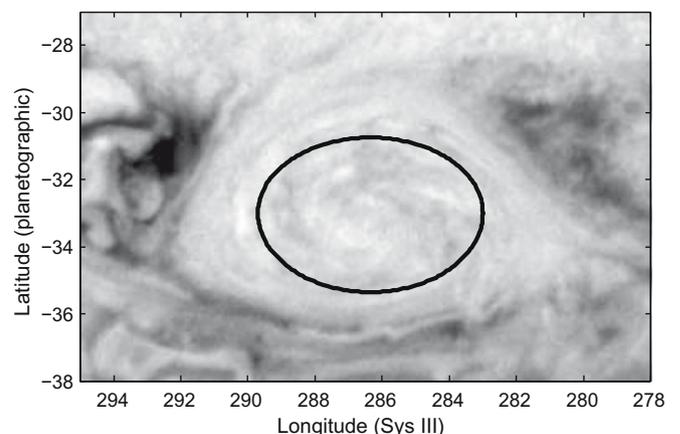
**Fig. 24.** Principal axis velocities of the Red Oval BA from HST images from April 2006 (data set *HST\_OBA\_06*). Velocity vectors were computed using CIV from images separated by about 40 min. Only points within  $0.2^\circ$  of the principal axes were used. The correlation uncertainty in the velocity vectors is  $\sim 9.5 \text{ m s}^{-1}$ . For comparison, Simon-Miller et al. (2006) found manual velocity vectors of the Oval using the same images we used to produce this figure and reported their uncertainties to be  $\sim 70 \text{ m s}^{-1}$ .



**Fig. 25.** Curves fit through the velocities on the (a) east-west and (b) north-south principal axes of the Oval BA. (c), (d), (e) and (f) close-up views of the velocity peaks, with the error bars indicating the averaged correlation uncertainties in the peak velocity magnitudes and locations. For the HST images taken in 2006, the correlation location uncertainty is not a useful measure of the uncertainty in the peak location due to the small separation time in the images. In most locations, the velocity field did not change between 2000 and 2006 to within the uncertainties. The exception is the east-west velocity at the northern edge of the vortex (panel e) due to the transient behavior discussed in Section 6.

through the center of the Oval BA shows that the velocities are unchanged between 2000 and 2006 at almost all locations. The uncertainty in the locations of the peak velocities is the RMS average of the magnitudes of all of the *correlation location uncertainties* of all of the independent tie-points within 500 km of each peak. For the 2006 velocity data, the RMS average of the *correlation location uncertainties* is 20 km, smaller than the effective resolution length and not a reliable indicator of the uncertainty of the peaks' locations (see Section 3.1).

Our results contradict those of Simon-Miller et al. (2006) who found a substantial,  $\sim 60 \text{ m s}^{-1}$ , increase in the maximum velocity of the Oval BA from 2000 to 2006. Simon-Miller et al. (2006) reported a large uncertainty,  $70 \text{ m s}^{-1}$ , in their velocities from 2006, based on the  $\sim 170 \text{ km}$  image pixel size of the original images (before deprojection) divided by the 40 min image separation time. This measure of uncertainty is not based on the velocity data themselves, and may be a significant underestimate of the true errors. The velocity field was composed of 236 independent vectors, which were extracted by hand from the same HST image pairs that we used in our CIV extraction.



**Fig. 26.** The cloud reflectivity of the Oval BA at 658 nm from the *HST\_OBA\_06* data set along with a black contour delineating the approximate outer boundary of the Oval BA (as in Fig. 21). The area of Oval's cloud cover is much larger than the area of the vortex.

Fig. 26 shows a smooth contour that approximates the local peak velocity along spokes radiating outward from the vortex center, superimposed on an image of the clouds of the Oval BA. Shetty (2008) showed this contour approximately corresponds to the outer edge of the potential vorticity anomaly of the Oval BA. Similar to the GRS, the Oval BA has a cloud cover with an area that is larger (in this case, much larger) than the area of the vortex itself.

Unfortunately, the Cassini images cannot be used for studying the long-term temporal changes of the northern part of the Oval BA because it was undergoing a transient interaction with the zonal jet on the northward side at the time Cassini flew by; the northern edge of the Oval BA was pulled north and distorted into a cusp-like shape with low velocity (Fig. 25f), and Shetty (2008) found that the Oval was far from equilibrium.

## 7. Errors and uncertainties

### 7.1. Navigation errors

In most cases, the images and mosaics that we used with ACCIV have documented navigation errors. The navigation errors of the *GLL\_GRS\_96* mosaics, *GLL\_GRS\_00* mosaics and the *CAS\_GRS\_00* and *CAS\_OBA\_00* images are  $\sim 20$  km (Vasavada et al., 1998),  $\sim 12$  km (Choi et al., 2007) and  $\sim 27$  km (Salyk et al., 2006), respectively. Our HST images (*HST\_GRS\_06* and *HST\_OBA\_06* data sets) had initial navigation errors of  $\sim 170$  km (see Appendix B.3). We were able to apply a correlation method similar to CIV to images less than  $\sim 40$  min apart in order to reduce the relative navigation error between images to a small fraction of a pixel. Since ACCIV requires the feature of interest to be placed in a steady-state frame (or one in which the characteristic time for the velocity to change is much greater than the separation time between images), we use the correlation method to find the frame closest to a steady state (a frame that may be translating with respect to System III). Despite the improved relative navigation of the HST images, the absolute navigation remains uncertain to  $\sim 110$  km. Fortunately, the absolute navigation error has no effect on the vorticity, and it only has the effect of placing the velocity field in a frame of reference that is translating with respect to the desired reference frame. For purposes of modeling, it is generally much more important to make sure that all images are in the same navigational reference frame than it is to know exactly what that reference frame is with respect to System III. Thus, the systematic navigational errors for all data sets are not important for modeling, and we found that we could either compensate for the relative navigational errors (in the case of the HST data sets) or that these errors were negligible (in the case of the Galileo and Cassini data sets).

### 7.2. Uncertainty measures used in previous work

Most manual techniques for tracking clouds have used the pixel size divided by the time interval between images to estimate the uncertainty in each velocity vector (Mitchell et al., 1981; Simon et al., 1998; Simon-Miller et al., 2006). This is based on the assumption that a user can find the location of a feature in each image to within a pixel. The uncertainty is independent of the actual data. Based on that estimate, the manual points in Fig. 3 would have an uncertainty of  $\sim 1.5$   $\text{ms}^{-1}$ , whereas our proposed correlation uncertainty (see Section 3.1) estimates a much larger uncertainty of  $\sim 8$   $\text{ms}^{-1}$ . What is more, ACCIV can find features with sub-pixel accuracy. This means that ACCIV (or CIV on its own) may find errors much smaller than would be estimated by the pixel size limit when images are separated by short times. Over long times other factors will almost always dominate the uncertainties in any technique.

The automated method of Choi et al. (2007) uses half the correlation box size multiplied by the horizontal shear of the flow field as their estimate of uncertainty, which we will refer to as the *shear uncertainty*. The uncertainty with this definition has little dependence on the actual data. The idea is that the dominant feature within the correlation box may be at the edge of the box rather than in the center. In this case, the resulting correlation shows the movement of the feature at the edge of the box rather than at the center of the box (as assumed). The velocity at the center differs from the velocity at the edge of the correlation box by the horizontal shear multiplied by the distance (half the box size). Based on this estimate, it would seem that one could reduce the error to nearly zero by shrinking the correlation box. In practice, there is a limit to how small the box can be before “false correlations” (see Section 7.3) begin to dominate the error. Choi et al. (2007) found reliable correlations using the *GLL\_GRS\_00* data set with a box size of 10 pixels, corresponding to a shear uncertainty of  $\sim 2$   $\text{ms}^{-1}$ . They suggested that the only other important source of error is image navigation. We find that many other factors (see Section 7.3) can dominate the error. Using ACCIV, we found a dramatic degradation of the velocity field when we changed the correlation box size from 30 pixels (the value used to extract ACCIV velocities from the *GLL\_GRS\_00* data set) to 10 pixels on a side. This is not surprising. The effective resolution of the *GLL\_GRS\_00* mosaics is 6–7 pixels, and we have found that the correlation box size must be a factor of 3–4 times the effective resolution length to produce reliable correlations.

ACCIV includes two features that remove distortions (and thus effectively eliminate the shear uncertainty): (1) CIV uses a correlation box that distorts with the flow (see Appendix A.1) and (2) ACCIV’s cloud advection step is specifically designed to “undo” cloud feature distortion. When we applied the shear uncertainty estimation method to the *Synth1* data set, we found an estimated uncertainty of 30% of the maximum velocity, which is 10 times the actual maximum error and 30 times the RMS error. We think that the shear uncertainty may significantly underestimate the error in velocities produced from images with short time separations (where other factors dominate), and would significantly overestimate the error in ACCIV velocities produced from images with long time separations since ACCIV compensates for the shear.

Del Genio et al. (2007) provided a measure of uncertainty that is based directly on their velocity data, as opposed to properties of their images or algorithm. Their velocity data, which covered a large portion of Saturn’s cloud deck, was produced using a 1-pass correlation technique that removed any correlation vectors with correlation coefficients less than 0.5. To measure the velocity uncertainty at a given latitude, they computed the difference between the mean value of *all* velocity measurements at that latitude and the mean value of the subset of velocity measurements at that latitude with correlation coefficients greater than 0.8. Using this measure, they found uncertainties of 2–10  $\text{ms}^{-1}$ . This uncertainty estimate may be a reliable estimate of the velocity error, though it has not been validated using synthetic data.

### 7.3. Causes of random errors in ACCIV

There are many sources of the random errors in velocities extracted from images of clouds. Clouds are created and dissipated, lighting conditions vary, images contain noise and Jupiter’s haze layer obscures clouds, especially near the limb. As long as these changes remain small, ACCIV will still be able to find correlations, but they will be slightly incorrect, leading to small random errors in the velocity. As the changes become more pronounced, ACCIV will no longer be able to find correlations in certain locations, leaving a “bald” spot in the velocity field. In the process of tracing the trajectories of neighboring features, ACCIV often patches over bald

spots. But sometimes there are too few neighbors and the bald spot persists. ACCIV may also find “false” correlations with nearby features that look similar. CIV uses several methods to remove false correlations (Fincham and Delerce, 2000). If false correlations are not removed they can create large random errors, but even if they are removed they leave bald patches behind.

Time dependence in the flow field also causes variations in the velocity vectors that we shall treat as errors. ACCIV assumes the velocity field to have negligible time dependence during the interval between images. Time dependence in the velocity field (or extracting velocities from clouds that are at different elevations in the atmosphere) can result in neighboring velocity vectors that are inconsistent, e.g., feature paths may cross. Fortunately, most of the time dependence in Jovian flows is due to turbulence with characteristic velocities much smaller than the time-independent dominant flow, which we expect to cause random, small amplitude variation in velocity vectors. One noteworthy exception is the turbulent wake to the northwest of the GRS, which appears to have significant time dependence – ACCIV often cannot find consistent paths for features in this region.

A final source of random errors is fundamental to CIV. CIV computes a small grid of correlation values as a function of displacement (the correlation function). A sub-pixel accurate displacement vector results from finding the maximum of the correlation function, using an interpolating function to compute values between grid points. Small random errors arise in interpolating the peak in the correlation function (Fincham and Delerce, 2000).

#### 7.4. Systematic errors

The methods for measuring uncertainties discussed in Sections 3.1 and 3.6 are designed to find only random errors; neither is able to take into account systematic errors inherent to ACCIV. Systematic errors show up in ACCIV velocity fields in two main ways. First, ACCIV systematically underestimates the magnitude of velocity vectors near peaks in the field. Second, CIV can often produce patches of false correlations which are removed, leaving behind bald spots, as discussed in Section 7.3.

Velocity peaks are rounded both because of the finite size of the CIV correlation box and because of limited *effective resolution* of the images (see Appendix B). With increasing box size or decreasing effective resolution, the velocity vectors measured by CIV will tend toward a local average value rather than the value at the center of the correlation box. The three Synthetic Gaussian Vortex data sets (*Synth1*, *Synth2* and *Synth3*) showed that ACCIV uncertainties are strongly dependent on effective image resolution. A more thorough sensitivity analysis is clearly warranted. (Note, this rounding is inherent to the velocity vectors produced by ACCIV and has nothing to do with the smoothing length for velocity interpolation discussed in Appendix A.2.)

Bald spots are another cause of systematic errors. They may be less of a problem than the rounding of velocity peaks depending on their size and position in the data set. The interpolation algorithm of Appendix A.2 will interpolate through bald spots based on neighboring data. The interpolated values will be nearly correct as long as the bald spot did not happen to occur at a location where variation in the velocity field was on a length scale smaller than the size of the bald spot. The most troublesome bald spots are those that occur at peaks in the velocity field, where the interpolated value of the velocity field will be lower than the actual velocity magnitude.

## 8. Conclusions

Historically, the clouds of the GRS, Oval BA, and other Jovian vortices have been used to characterize the underlying vortices.

However, we have shown here that those clouds are not aligned with any dynamical features of the vortices, and their areas are greater than those of the actual vortices. To understand the dynamics of Jovian vortices it is necessary to compute their velocity fields and not just track the morphological changes in their clouds. Using the ACCIV and CIV automated methods for extracting velocity fields from image pairs, we have shown that, to within our reported uncertainties, the velocities and areas of the Oval BA did not change from December 2000, when it was white, to April 2006, when it was red. Using ACCIV, our new velocity-extraction algorithm, the velocities of the Oval BA have *correlation* and  $1\sigma$  *uncertainties* as low as  $6.5 \text{ m s}^{-1}$ , and  $3 \text{ m s}^{-1}$ , respectively. Using CIV, an established automated method for extracting velocities, the *correlation* and  $1\sigma$  *uncertainties* are as low as  $9.5 \text{ m s}^{-1}$  and  $5.5 \text{ m s}^{-1}$ , respectively, which is remarkable considering that the time separation between the images used in the extraction is only 40 min. We determined that the cloud lifetimes of the Oval BA were greater than 10 h in 2000, but less than 10 h in 2006.

Using ACCIV, we have shown that the area circumscribed by the high-speed collar of the GRS (which is approximately coincident with the potential vorticity anomaly of the GRS (Shetty et al., 2007; Shetty, 2008)) decreased by 15% from June 1996 to April 2006. The change in area is real and not an observational artifact; the change in area was significantly greater than the uncertainties in our measurements. During the same interval of time, the GRS became rounder, but its peak velocity magnitudes remained constant (to within the *correlation* and  $1\sigma$  *uncertainties*, which are as small as  $2 \text{ m s}^{-1}$  and  $1.2 \text{ m s}^{-1}$ , respectively). There are several scenarios that can account for the change in areas and shape of the GRS (e.g., the stripping of potential vorticity as the GRS encounters stagnation points, change in the potential vorticity gradients of the ambient zonal flows, a decrease in the rate at which the GRS merges with small vortices, changes in the vertical stratification, etc.), and these scenarios have been modeled by Shetty (2008). Although the results of those analyses are beyond the scope of this paper, it is important to note that the modeling could not have been initialized and the output of the models could not have been validated without the use of highly-accurate velocity fields with tens to hundreds of thousands of independent vectors with well-quantified uncertainties. Given the quality of archived spacecraft images, fields of this kind can only be obtained with ACCIV. The reason that ACCIV can produce velocities with uncertainties that are much smaller than other automated methods is that ACCIV can find correlations in images pairs that are separated by 10 h, whereas other automated methods cannot.

Based on recent publications (Table 1), manual, rather than automated, methods remain the preferred way of extracting velocities from high resolution spacecraft images of Jupiter. This is surprising, because most fluid dynamicists studying laboratory flows abandoned manual methods for automated ones more than 15 years ago (Sommeria et al., 1988, 1989, 1991). It was realized that manually-extracted velocity fields were too sparse to be useful in most analyses. The reluctance of planetary scientists to embrace automated methods may be due to misconceptions about the relative merits of the two methods. To confront these misconceptions, we carried out the first objective comparison of automated and manual methods by applying both methods to synthetically advected clouds. The results of those test showed that the “averaging” of velocities with ACCIV is not large. The rounding of velocity peaks due to the inherent smoothing and averaging in ACCIV was found to be *less than* the rounding of the peaks of the manually-extracted velocity peaks. The rounding in manual methods is due to the approximation of curved cloud trajectories with straight lines. This approximation is only valid when the time separation between images is small (or if one confines studies to zonal flows without vortices), though the small time separation will lead to

large random errors. Our analyses with synthetic clouds confirmed that ACCIV achieves sub-pixel accuracies in the locations of the tie-points used to calculate the velocity vectors, but surprisingly, the manually-extracted tie-points had errors in location that were significantly greater than one pixel. Although the magnitudes of the peak velocities found by ACCIV were slightly smaller than the exact values, ACCIV found the locations of the peaks to sub-pixel accuracy. On the other hand, the magnitudes of the peak velocities found with manual methods were even smaller than those found by ACCIV; moreover, the locations of the peak velocities of the synthetic vortices that were found by manual methods were systematically skewed toward the vortex center, so that the area of the manually-extracted vortex appeared smaller than its exact value.

We have defined two new measures of the uncertainty of extracted velocity fields. These new uncertainty estimates more accurately reflect the errors than the two commonly-used current theoretical uncertainty estimates: (1) pixel size divided by the separation time between images (Mitchell et al., 1981; Sada et al., 1996; Simon et al., 1998; Simon-Miller et al., 2006); and (2) the *shear uncertainty* (Choi et al., 2007). These commonly-used estimates of the uncertainty are unrealistic because they are independent of image quality and give the same results regardless of whether the images are blurry or noisy, whether the cloud feature contrast is high or low, or whether the velocity field has significant turbulence. Our new measures of uncertainty include these effects. In future work, it would be useful to carry out a full analysis of the sensitivity of our uncertainty measures (which can be applied to any velocity field, regardless to how it is extracted) to the user-specified parameters of the extraction method, e.g., correlation box size, smoothing length, number of vectors along a cloud path, etc. For example, using synthetic data, we showed that the ACCIV uncertainties are sensitive to effective image resolution (see Fig. 9) and have a much weaker dependence on the pixel resolution of the image. We also showed that the ACCIV velocity is nearly independent of both the smoothing length used in tie-point trajectories and the correlation box size. Much more analysis could be done. The sensitivity analysis could give us a better grasp of the systematic errors in ACCIV that are not included in our current uncertainty estimates. It may even be possible to compensate for errors caused by limited image resolution and finite correlation box size by using the sensitivity analysis results to extrapolate to the limits of infinite image resolution and zero box size.

In addition to the GRS and Oval BA, we have applied ACCIV to regions of our HST Jovian image pairs from April 2006 that do not contain any long-lived vortices, and the method successfully produced hundreds of thousands of independent velocity vectors with small uncertainties. This means that ACCIV can be used to map nearly the entire planet and search for waves, hidden vortices and other phenomena.

## Acknowledgments

The new observations presented here of Jovian cloud images were made with the NASA/ESA Hubble Space Telescope. These observations are associated with HST program GO/DD 10782, with support provided by NASA through a grant from the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555. Analysis was supported by grants from the Planetary Atmospheres Program of NASA and the Astronomy and Astrophysics Program of NSF, with a computational allocation from the TeraGrid funded by NSF.

We thank Ashwin Vasavada and David Choi for sharing processed Galileo and Cassini imaging data, Reta Beebe for sharing Voyager 1 velocity data, Sean Lockwood for helping with HST data processing and software development, and to Christopher Go for

inspiring and motivating this work by sharing his observations of Jupiter, including his discovery of the Oval BA's color change. Thanks also to Ashwin Vasavada for a thorough and collegial review of this work.

## Appendix A. Details of the ACCIV algorithm

The steps in the ACCIV algorithm to extract velocities from images with a 10-h separation in time are:

- (1) Produce an initial set of discrete velocity vectors (at locations, that in general, are not on a uniform grid). To do this, either apply CIV (or some other automated method) to images separated by a short time interval (2 h or less) or find manual correlations over a 10-h interval.
- (2) Using these discrete vectors, interpolate the velocity onto a uniform grid to create an initial estimate of the velocity *field*.
- (3) Using the grid velocity, numerically advect the clouds by moving the pixels in each image forward or backward to a single time. For example, if the first image showing the clouds is at time  $t = 0$  h and the second image is at  $t = 10$  h, then advect the first image forward to time  $t = 5$  h and the second image backward in time to  $t = 5$  h. Thus, there will be different synthetic images at  $t = 5$  h. Often, the image  $t = 0$  h will have a different estimated velocity field than the image at  $t = 10$  h, since each velocity field was produced by applying CIV to a different subset of images.
- (4) Use CIV to identify correlations between the clouds in the two *advected* images (i.e., at  $t = 5$  h). Normally, we would identify the correlations as clouds displacements, but in this case, the correlations are “correction vectors” or *errors* in the estimated velocity field (or fields) that was used to advect the clouds. For example, if the velocity field were perfect, CIV would find millions of correlations and the distances between the tie-point pairs would be zero. If the estimated velocity field were reasonably good, CIV would still find millions of correlations, but the distances between tie-points would be large. If the estimated velocity field were very poor, it is possible that CIV would find very few or no correlations.
- (5) Trace each feature forward or backward in time to find its location in the original images, essentially undoing the advection process.
- (6) Produce a new set of discrete velocity vectors and a new velocity field on a uniform grid from the cloud feature displacements by iteratively tracing streamlines of the flow.

Steps 3 through 6 are iterated until no further improvement is seen in the velocity field. Typically, this requires  $\sim 3$  iterations.

We now describe these steps in detail.

### A.1. Estimating initial velocities

ACCIV relies on having a prediction of the velocity field that can be iteratively corrected to improve accuracy, sometimes by more than factor of six (e.g.,  $32 \text{ m s}^{-1}$  uncertainty for the CIV velocity data compared with  $5 \text{ m s}^{-1}$  uncertainty for the ACCIV velocity data produced from the *HST\_GRS\_06* data set). We have found that applying CIV to images separated by 2 h or less is the simplest way to get a rough estimate of the velocity field. Alternatively, we could find the predicted velocity field by applying another automated method or by manually correlating cloud features over  $\sim 10$  h. The initial estimated velocities do not need to be accurate; we have found that a field with an uncertainty as large as  $32 \text{ m s}^{-1}$ , about 30% of the peak velocity, can still be used as the first guess.

CIV works in two passes, shown in Fig. 27. The first pass involves translating a correlation box of pixels in an earlier image to find the best-correlated matching box of pixels in a second image. The second pass makes use of a smaller correlation box that is both translated and deformed based on the information provided by the first pass. The smaller correlation box size means that correlations will have higher spatial resolution, while deformation of the correlation box increases correlation accuracy and reduces so-called “peak locking errors”, which limit the sub-pixel accuracy of most other automated methods.

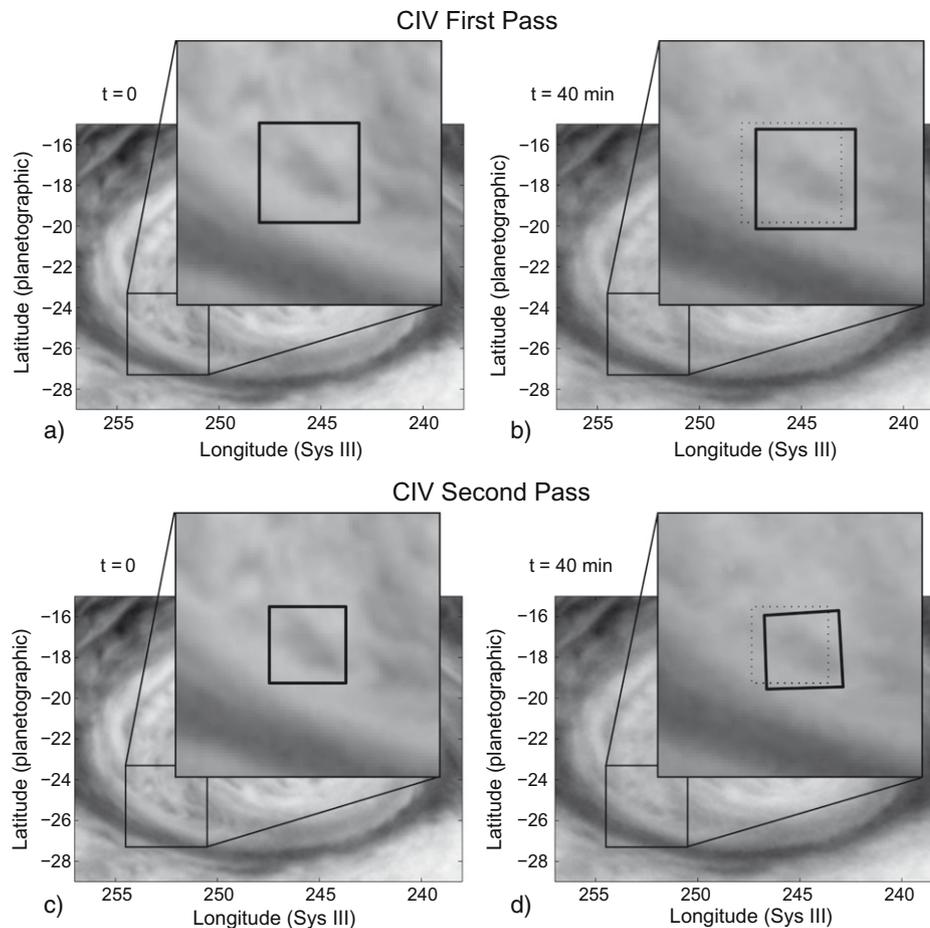
#### A.2. Interpolating the velocity to and from a uniform grid

Image advection requires knowing the velocity at all locations – not just at the locations of the initial discrete velocity vectors and not just at the grid points. This is because each pixel in the images needs to be numerically advected. At most time steps during the advection, the pixel is located neither at the location of a discrete velocity vector nor at a grid point. This means that we need a velocity field on a regular grid so that the velocity at an arbitrary location can be found quickly by interpolation. The “thin plate spline” method (Sandwell, 1987) has commonly been used in laboratory flow experiments (Sommeria et al., 1988, 1989) to produce a velocity field on a grid from a small number of velocity vectors. The velocity field passes through all data points, which leads to a

very jagged field when the data is noisy and dense. Also, the number of computations required by common thin plate spline algorithms scales as  $n^3$ , where  $n$  is the number of data points. This means that thin plate splines are ill suited to the numerous, dense and noisy data that CIV and ACCIV produce.

Instead, we use a smooth surface fitting technique, *gridfit*, developed for Matlab by D’Errico (2006). The smooth velocity field is the least-squares solution of the over-determined linear system requiring the smooth function to match the data and the function’s second derivative in each direction to be zero. Values of the smooth function at the data points are found by linear interpolation, and derivatives are computed by second-order finite differences. The least-squares solution will not be able to satisfy all the constraints (that is what makes the system over-determined) but the solution will come “as close as possible” to satisfying them. In practice, this leads to function that is both smooth (its second derivatives are small) and a good fit to the data.

The fitting technique requires a user-specified constant, the *smoothing length*, to determine the relative weight of smoothness and closeness of fit. We have not explored automated methods for determining the smoothing length although this may warrant future study. Instead, we fit fields to the data based on a series of test smoothing length values. We manually chose the “best” smoothing length so that noise is smoothed without losing significant flow features. Fortunately, the iterative nature of ACCIV



**Fig. 27.** Correlation boxes in the first pass (a) and (b) and second pass (c and d) of the CIV algorithm. (a) The inset shows the correlation box of pixels from the first image at  $t = 0$  as a solid line. (b) The inset shows the second image at  $t = 40$  along with the same *untranslated* correlation box from (a) (dotted line) and the *translated* correlation box, corresponding to  $t = 40$  (solid line). (c) In the second CIV pass, a correlation box of pixels from the first image (usually smaller than the correlation box in the first pass) at  $t = 0$  is shown with a solid line in the inset. (d) In the inset, the location in the second image of that same and *untranslated and undistorted* correlation box as in (c) (dotted line) and the *translated and distorted* location of that box at  $t = 40$  (solid line).

means that the exact amount of smoothing during the intermediate steps generally does not have a large effect on the final velocity field (Fig. 14). During the final ACCIV iteration, it may be more important to choose the smoothing length with care. Primarily, this is because we use the difference between the smooth velocity field and the individual velocity vectors to compute the  $1\sigma$  uncertainty. Also, the smooth velocity field will be used to determine the curvature of feature paths, which will, in turn, influence the final velocity vectors (see Appendix A.6).

### A.3. Simulating cloud advection

With the gridded velocity field, we advect the pixels of each image either forward or backward to a the point half-way in time between the first and second images. The gridded velocity field allows us to efficiently interpolate the velocity at arbitrary points in a way that would be prohibitively expensive if we were attempting to use the original scattered data. We use a spectral method that is second-order accurate in time to advect the images. Any reasonably high accuracy method for advecting passive tracer fields should suffice.

### A.4. Identifying cloud features in the advected images

If the velocity field were exactly correct, and the clouds changed only by moving with the flow, the advected images at the half-way time would exactly match one another. Of course, in practice the velocity field is not free of errors, and the cloud images change due to a number of processes other than advection, so that the advected images will never match exactly. Applying CIV between the advected images results in a large number of correction vectors that show how misaligned the cloud features are. The correction vectors can be used to estimate the local and global uncertainties of the velocity field, as discussed in detail in Section 3.1.

### A.5. Tracking features back to the original images

The correction vectors have no direct physical meaning; they do not represent the displacement of cloud features along a path as they did in the first use of CIV. To get the real displacement vector and the real tie-point locations, we need to “undo” the advection process. That is, we need to track the correlated features in the first image that was advected to the half-way time backward in time 5 h to their true tie-point locations in the unadvected first image. We also track the features in the second advected image at the half-way time forward 5 h to their places (true tie-points) in the later unadvected image. Fig. 28 shows four circled features being tracked from the advected images (at the top) back to the unadvected images (at the bottom). With this method, we have indirectly found tie-points between the first and second images that CIV was unable to find directly.

### A.6. Finding curved cloud paths to produce velocity vectors

Displacement vectors are not sufficient to produce velocity vectors if cloud paths are curved, as is frequently the case over 10 h. The location, magnitude and direction of the velocity will be significantly misrepresented if the velocity vector is constructed in the usual way, that is, by taking the displacement divided by the time interval and placing the vector at the midpoint between feature locations. A better approach is to trace the cloud path, placing a sequence of velocity vectors along it. Fig. 29 shows how big a difference this approach can make. In particular, the velocity of the feature on the left would be completely misrepresented by treating its displacement (divided by the time interval) as its velocity.

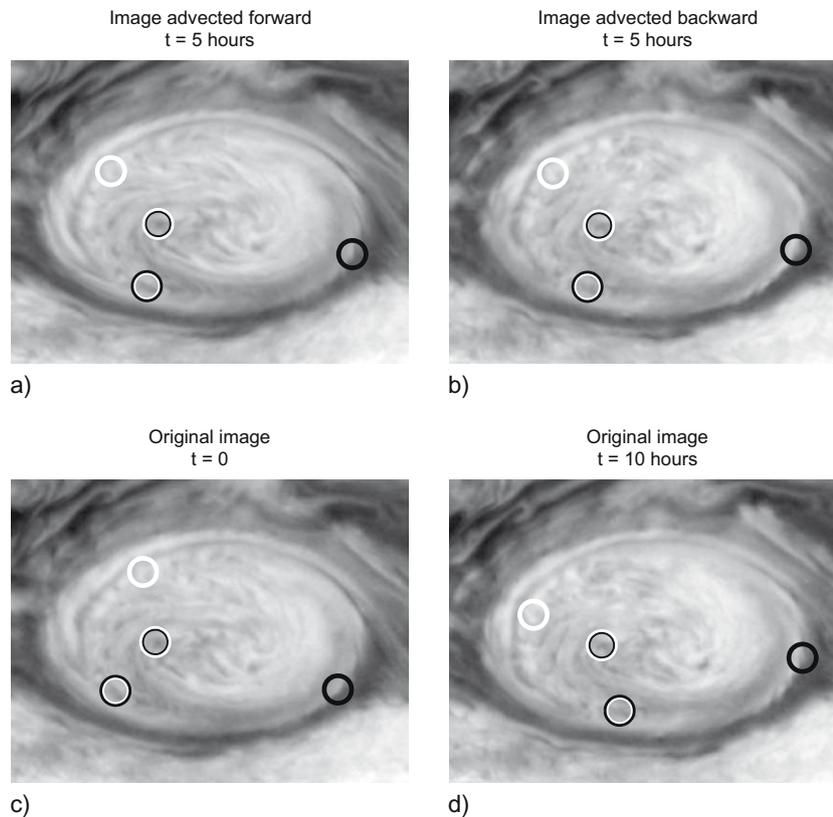
To produce accurate velocity vectors, we need a way to derive cloud paths from displacements. We accomplish this iteratively. As a starting point, we assume no curvature and produce a single velocity vector from each displacement by dividing by the separation time. Then, we produce a velocity field on a uniform grid using the same surface fitting technique described in Appendix A.2. With the new velocity field, we advect each feature forward from the first image to the second image, and also backward from the second image to the first. Since the two paths do not agree exactly, we linearly interpolate between the two trajectories, so that the resulting path takes the feature smoothly between its known locations in the two images (see Fig. 8). To do this, let  $(\bar{X}_1(t), \bar{Y}_1(t))$  (upper dashed path) be the trajectory that leads from the actual (solid circle) location of the tie-point in the first image,  $(x(t_1), y(t_1))$  to its erroneous location (open circle) in the second image,  $(x'(t_2), y'(t_2))$ , as computed numerically by forward integration from  $t_1$  to  $t_2$  with the assumed velocity field. Let  $(\bar{X}_2(t), \bar{Y}_2(t))$  (lower dashed path) be the trajectory that leads back from the actual location (solid circle) of the tie-point in the second image,  $(x(t_2), y(t_2))$  to its erroneous location (open circle) in the first image,  $(x'(t_1), y'(t_1))$ , as computed numerically by backward integration from  $t_2$  to  $t_1$  with the assumed velocity field. The linearly interpolated path is defined as  $(\bar{X}(t), \bar{Y}(t)) \equiv [(t_2 - t)/(t_2 - t_1)](\bar{X}_1(t), \bar{Y}_1(t)) + [(t - t_1)/(t_2 - t_1)](\bar{X}_2(t), \bar{Y}_2(t))$ . We create new, artificial velocity vectors (typically, 4–8) along the interpolated path by taking the local time derivative of  $(\bar{X}(t), \bar{Y}(t))$ . Using the set of artificial velocity vectors, we compute a new velocity field at the grid points. We iterate this process until the cloud feature paths do not change significantly from one iteration to the next.

### A.7. Iterating ACCIV

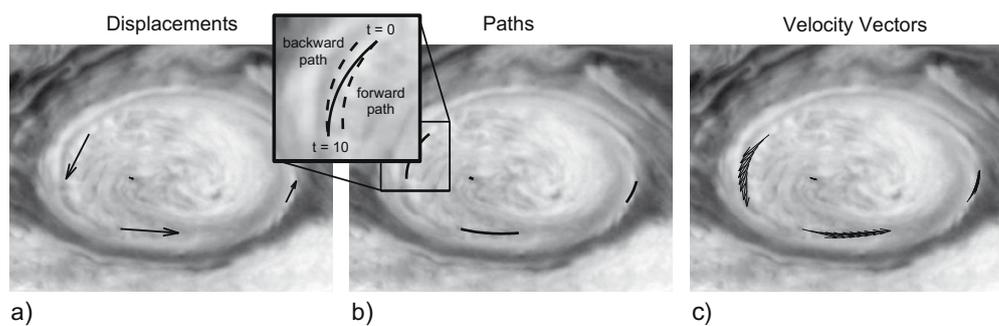
We iterate steps 3 to 6 – advecting the images using the updated velocity field, finding new correction vectors, and producing a new set of velocity vectors and a new velocity field. In a small number of iterations (usually about three), the velocity field ceases to improve with iteration. The rate of convergence of these steps can be increased slightly by removing outliers from the set of velocity vectors. We do this by first interpolating the velocity field onto the grid from the full set of vectors at their given (non-grid) location. Then, we create a new velocity at the given (non-grid) locations of the original set of vectors by interpolating the velocity from the grid back to those given (non-grid) locations. We eliminate outliers by removing all velocity vectors whose original values and new values at the given (non-grid) locations differ by more than six times the median value of those differences. Then, a new velocity at the grid points is constructed using the remaining velocity vectors. Inevitably, some of the velocity vectors removed by this technique could be accurate, so we only perform this process during intermediate iterations of ACCIV. For the extracted velocity fields presented in this paper, outliers were not removed in the last iteration of ACCIV.

### A.8. Applying ACCIV to the SynthCusp data set

To give the reader a better idea of the ACCIV process, we will describe in detail its application to the *SynthCusp* data set. As the first step in ACCIV, we used CIV to extract two velocity field estimates, one each from the four earlier images and one from the four later images. The first CIV pass used a correlation box length of 20 pixels and with a search range of  $\pm 20$  pixels, while the second CIV pass used a correlation box length of 15 pixels. The search range of the second pass is determined automatically by the CIV algorithm, and is confined to be within a few pixels of the results of the first CIV pass. Using the 6 possible pairings among the 4 earlier images



**Fig. 28.** The “undoing” of the advection step in the ACCIV algorithm. (a) An image created in an ACCIV subroutine by *advecting* a real image at  $t = 0$  forward 5 h in time to the half-way time at  $t = 5$  h. (b) An image created by ACCIV by *advecting* a real image at  $t = 10$  backward 5 h in time to the same half-way time at  $t = 5$  h. ACCIV uses CIV to find correlations of the advected features in (a) and (b). Four features are circled in each image with differently highlighted circles. (c and d) The advection process described above is then reversed, starting with images (a) and (b) at the half-way time, to find the locations of the features in the original (real) images at  $t = 0$  and  $t = 10$  h. The real displacement vectors point from the features’ locations in (c) to their locations in (d).



**Fig. 29.** Steps taken by ACCIV to produce artificial velocity vectors from feature displacements. (a) Vectors showing the displacement of the four features from Fig. 28c to Fig. 28d. (b) The inset shows two paths with dashed lines: one is the path of one of the cloud feature as it is advected forward 10 h in time from its location in the first image ( $t = 0$ ), and the second is the path of that same feature as it is advected backward 10 h in time from the second image ( $t = 10$ ) (see Fig. 8). The two paths are not coincident, so linear interpolation between the paths is used to find a “better” path (solid line). (c) Artificial velocity vectors are placed along the interpolated paths.

in this data set, CIV produced 294,689 correlation vectors, of which we estimate  $\sim 150,000$  were independent. Curved paths were constructed for each of these vectors, and 4 artificial velocity vectors were placed along each path for a total of 1,178,756 (non-independent) velocity vectors. After removing outliers (see Appendix A.7), 1,094,534 velocity vectors remain. From these vectors, we construct a velocity field on a regular grid using a smoothing length of 4 pixels to produce a velocity field on a grid. Similar numbers of vectors were produced from the 4 later images in the data set; a second estimated velocity field was produced from these data. Using the estimated velocity fields on the grid, the ACCIV algorithm

advected the synthetic images to  $t = 0.110$ , about half-way between the earlier and later images. We found corrections between the two advected images using CIV, where the first pass had a correlation box length of 40 pixels and the correlation box length in the second pass was 20 pixels. The search range for the first pass was  $\pm 20$  pixels. We produced displacements from the correction vectors by advecting the heads forward in time and the tails backward in time to the times of the original images. Then, ACCIV iteratively produced curved paths between the feature locations as described in Section 3.3 and as shown in Fig. 8. New, artificial velocity vectors were placed along the paths. As before, ACCIV used

a smoothing length of 4 pixels. Finally, ACCIV advected the original images back to  $t = 0.110$  using the new velocity field. We repeated the entire ACCIV algorithm twice more, each time using the updated velocity field on the grid to advect the images. During the third iteration, CIV produced 677,152 correction vectors, of which  $\sim 350,000$  were independent. The displacement vectors were used to make 2,708,608 non-independent velocity vectors by placing 4 vectors along each feature path. We did not remove outliers in the final iteration.

## Appendix B. Observations used for velocity extraction with ACCIV for the GRS and Oval BA

### B.1. Galileo Solid State Imaging (SSI) data

We used two sets of Galileo mosaics of the GRS. The first set of three mosaics from June 26, 1996 was provided by Ashwin Vasavada of the Jet Propulsion Laboratory, California Institute of Technology, who processed them using methods described in Vasavada et al. (1998). The deprojected mosaics are available through the Atmospheres Node of the Planetary Data System (PDS),<sup>1</sup> where they have been identified as

G1JSGRSEM101-large,  
G1JSGRSEM201-large,  
G1JSGRSEM301-large.

In this set the second and third mosaics were taken 9 and 10 h, respectively, after the first mosaic. The images were taken at 756 nm with a resolution of  $\sim 30\text{--}36$  km  $\text{pix}^{-1}$ . The deprojected mosaics have oversampled the data to a resolution of  $\sim 18$  km  $\text{pix}^{-1}$ . We find that the effective image resolution is  $\sim 52$  km. Throughout the paper, these images are referred to as the *GLL\_GRS\_96* data set.

The second set of three mosaics (data set *GLL\_GRS\_00*) from May 2000 were provided by David S. Choi of the University of Arizona, and are described in Choi et al. (2007). Each mosaic is separated from the previous by about 1 h. Similar (but not identical) versions of these mosaics are available through the PDS Atmospheres Node under the names:

G28JSFEATRK01-large,  
G28JSFEATRK02-large,  
G28JSFEATRK03-large.

These mosaics were imaged at 756 nm wavelength with an original image resolution of  $\sim 17$  km  $\text{pix}^{-1}$ . The deprojected mosaics are oversampled to  $\sim 11$  km  $\text{pix}^{-1}$  resolution, and have an effective resolution of  $\sim 71$  km. Note that, though the pixel resolution of these images was higher than the Galileo images from 1996; the effective resolution was lower. The lower effective resolution is probably related to considerable amount of noise in the 2000 mosaics compared with those 1996. As far as we are aware, the source of this additional noise is unknown.

### B.2. Cassini Imaging Science Subsystem – Narrow Angle (ISS-NA) data

We use 8 images of the GRS (data set *CAS\_GRS\_00*) and 8 images of the Oval BA (data set *CAS\_OBA\_00*) taken from a sequence of 57 Cassini image maps from December 11–13, 2000. These image maps were provided by Ashwin Vasavada, who processed these images as described in Porco (2003) and Vasavada et al. (2006). These images were acquired at 750 nm with  $\sim 114\text{--}122$  km  $\text{pix}^{-1}$

resolution. The images were deprojected with comparable resolution of  $\sim 110$  km  $\text{pix}^{-1}$ , and  $\sim 270$  km effective resolution at the latitude of the GRS and Oval BA. The cylindrically deprojected images containing the GRS are available from the PDS Atmospheres Node (see above) under the names:

n1355286849.1,  
n1355290635.1,  
n1355294421.1,  
n1355298207.1,  
n1355320923.1,  
n1355324709.1,  
n1355328495.1,  
n1355332281.1.

The comparable PDS images containing the Oval BA are:

n1355305779.1,  
n1355309565.1,  
n1355313351.1,  
n1355317137.1,  
n1355339853.1,  
n1355343639.1,  
n1355347598.1,  
n1355351384.3.

### B.3. HST Advanced Camera for Surveys (ACS) data

In April 2006, we acquired images of the GRS and Oval BA using the High Resolution Channel (HRC) of HST's Advanced Camera for Surveys (ACS). Data were photometrically calibrated using the inverse sensitivity parameter PHOTFLAM generated by the ACS calibration pipeline (Pavlovsky et al., 2006). New values of PHOTFLAM differing by less than 1% became available in 2007, but because our velocity retrievals are sensitive to relative rather than absolute reflectivity, we did not recalibrate the data. Data were converted to *I/F* using Eq. 1 of Stromovsky and Fry (2002). Cosmic ray strikes and geometric distortion were removed using *crutil* and *drizzle* within *PyRAF*.<sup>2</sup>

The processed images were deprojected into latitude–longitude space by fitting the limb of Jupiter where available, or by matching multiple cloud features to other limb-fitted images where no limb was available within the image. We created deprojection software that allows a human operator to align a synthetic Jupiter image against the real Jupiter data with quarter-pixel accuracy. The synthetic image includes Lambertian limb darkening and is convolved with the appropriate HST point spread function generated by the *Tiny Tim* software (Krist and Hook, 1997). Pointing uncertainty is estimated by computing the standard deviation of three navigation solutions: the user-generated solution, the solution from maximizing the cross-correlation between the limbs of the synthetic and real Jupiter images, and the solution from minimizing the standard deviation of the ratio of the synthetic and real Jupiter images (also at the limb). For images navigated using tie-points rather than limb alignment, navigation uncertainty was again estimated as the standard deviation of the same three solutions for each tie-point, and five tie-points per image were used. Based on this method, the navigation uncertainties were  $\sim 3$  pixels ( $\sim 170$  km). HRC exposures took advantage of HST's capability to track rotating planetographic

<sup>1</sup> <http://pds-atmospheres.nmsu.edu/Jupiter/jupiter.html>.

<sup>2</sup> *Drizzle*, *MultiDrizzle*, and *PyRAF* are products of the Space Telescope Science Institute, which is operated by AURA for NASA. *Crutil* is a package within *IRAF* (Tody, 1993). Because multiple images of a rapidly rotating spheroid cannot be constructively combined, the added features of *MultiDrizzle* cannot be used for mosaic composition and cosmic ray cleaning of planetary images.

coordinates on Jupiter's 1-bar surface, minimizing blurring due to planetary rotation.

ACCIV worked most successfully on images taken at 658 nm wavelength since these had the highest cloud feature contrast. The images before deprojection had a resolution of  $\sim 170 \text{ km pix}^{-1}$  at the equator. The deprojected images were over-sampled onto a grid with  $\sim 55 \text{ km pix}^{-1}$  resolution and with an effective resolution of  $\sim 163 \text{ km}$ . The original images (before deprojection) are available through the STScI MAST Archives,<sup>3</sup> where they are identified as:

J9MM04BPQ,  
J9MM04C4Q,  
J9MM06ECQ,  
J9MM06ELQ,  
J9MM04C7Q,  
J9MM06EFQ,  
J9MM06EOQ,  
J9MM04BMQ.

These images were cropped to the regions around the GRS for the *HST\_GRS\_06* data set and to the region of the Oval BA for the *HST\_OBA\_06* data set. Upon publication of this article, our deprojected HST/ACS maps will be uploaded as High Level Data Products to the HST archive, where they can be located using the identifiers above.

## References

- Cheng, A.F., Simon-Miller, A.A., Weaver, H.A., Baines, K.H., Orton, G.S., Yamamandra-Fisher, P.A., Mousis, O., Pantin, E., Vanzi, L., Fletcher, L.N., Spencer, J.R., Stern, S.A., Clarke, J.T., Mutchler, M.J., Noll, K.S., 2008. Changing characteristics of Jupiter's Little Red Spot. *Astron. J.* 135, 2446–2452.
- Choi, D., Banfield, D., Gierasch, P., Showman, A., 2007. Velocity and vorticity measurements of Jupiter's Great Red Spot using automated cloud feature tracking. *Icarus* 188, 35–46.
- de Pater, I., Marichis, F., Macintosh, B., Roe, H., Mignant, D.L., Graham, J., Davies, A., 2004. Keck AO observations of IO in and out of eclipse. *Icarus* 169, 250–263.
- Del Genio, A.D., Barbara, J.M., Ferrier, J., Ingersoll, A.P., West, R.A., Vasavada, A.R., Spitale, J., Porco, C.C., 2007. Saturn eddy momentum fluxes and convection: First estimates from Cassini images. *Icarus* 189, 479–492.
- D'Errico, J., 2006. Gridfit. Available from: <<http://www.mathworks.com/matlabcentral/fileexchange>>.
- Dowling, T.E., Ingersoll, A.P., 1988. Potential vorticity and layer thickness variations in the flow around Jupiter's Great Red Spot and White Oval BC. *J. Atmos. Sci.* 45, 1380–1396.
- Fincham, A., Delerce, G., 2000. Advanced optimization of correlation imaging velocimetry algorithms. *Exp. Fluids* 29, S013–S022.
- Fincham, A.M., Spedding, G.R., 1997. Low cost, high resolution DPIV for measurement of turbulent fluid flow. *Exp. Fluids* 23, 449–462.
- Gierasch, P.J., Ingersoll, A.P., Banfield, D., Ewald, S.P., Helfenstein, P., Simon-Miller, A., Vasavada, A., Breneman, H.H., Senske, D.A. Galileo Imaging Team, 2000. Observation of moist convection in Jupiter's atmosphere. *Nature* 403, 628–630.
- Krist, J., Hook, R., 1997. *Tiny Tim User's Manual*, v. 44. Space Telescope Science Institute, Baltimore, MD.
- Mitchell, J.L., Beebe, R.F., Ingersoll, A.P., Garneau, G.W., 1981. Flow fields within Jupiter's Great Red Spot and White Oval BC. *J. Geophys. Res.* 86 (A10), 8751–8757.
- Pavlovsky, C., Koekemoer, A., Mack, J., 2006. ACS Data Handbook, Version 5.0. Space Telescope Science Institute, Baltimore, MD.
- Perlin, K., 2002. Improving noise. *ACM Trans. Graph.* 21 (3), 681–682.
- Porco, C.C., and 23 colleagues, 2003. Cassini imaging of Jupiter's atmosphere, satellites, and rings. *Science* 299, 1541–1547.
- Read, P., Gierasch, P., Conrath, B., Yamazaki, Y., 2005. 3d balanced winds and dynamics in Jupiter's atmosphere from combined imaging and infrared observations. *Adv. Space Res.* 36, 2187–2193.
- Rossow, W.B., del Genio, A.D., Eichler, T., 1990. Cloud-tracked winds from Pioneer Venus OCPP images. *J. Atmos. Sci.* 47, 2053–2084.
- Sada, P.V., Beebe, R.F., Conrath, B.J., 1996. Comparison of the structure and dynamics of Jupiter's Great Red Spot between the Voyager 1 and 2 encounters. *Icarus* 119, 311–335.
- Salyk, C., Ingersoll, A.P., Lorre, J., Vasavada, A., Genio, A.D.D., 2006. Interaction between eddies and mean flow in Jupiter's atmosphere: Analysis of Cassini imaging data. *Icarus* 185, 430–442.
- Sandwell, D.T., 1987. Biharmonic spline interpolation of geos-3 and seasat altimeter data. *Geophys. Res. Lett.* 2, 139–142.
- Sheehy, C.D., McCrady, N., Graham, J.R., 2006. Constraining the adaptive optics point-spread function in crowded fields: Measuring photometric aperture corrections. *Astrophys. J.* 647, 1517–1530.
- Shetty, S., 2008. Formulation and Solution of the Inverse Problem for Jupiter's Atmospheric Vortices. Ph.D. Thesis, University of California, Berkeley.
- Shetty, S., Asay-Davis, X.S., Marcus, P.S., 2007. On the interaction of Jupiter's Great Red Spot and zonal jet streams. *J. Atmos. Sci.* 64, 4432–4444.
- Simon, A.A., Beebe, R.F., Gierasch, P.J., Vasavada, A.R., Belton, M.J.S. the Galileo Imaging Team, 1998. Global context of the Galileo-E6 observations of Jupiter's White Ovals. *Icarus* 135, 220–229.
- Simon-Miller, A.A., Chanover, N.J., Orton, G.S., Sussman, M., Tsavaris, I.G., Karkoschka, E., 2006. Jupiter's white oval turns red. *Icarus* 185, 558–562.
- Simon-Miller, A.A., Gierasch, P.J., Beebe, R.F., Conrath, B., Flasar, M., Achterberg, R.K. the Cassini CIRS Team, 2002. New observational results concerning Jupiter's Great Red Spot. *Icarus* 158, 249–266.
- Sommeria, J., Meyers, S.D., Swinney, H.L., 1988. Laboratory simulation of Jupiter's Great Red Spot. *Nature* 331, 689–693.
- Sommeria, J., Meyers, S.D., Swinney, H.L., 1989. Laboratory model of a planetary eastward jet. *Nature* 337 (6202), 58–61.
- Sommeria, J., Meyers, S., Swinney, H., 1991. Experiments on vortices and Rossby waves in eastward and westward jets. In: Osborne, A. (Ed.), *Nonlinear Topics in Ocean Physics*. Name of Publishing Company, North-Holland, Amsterdam, pp. 227–269.
- Sromovsky, L., Fry, P., 2002. Jupiter's cloud structure as constrained by Galileo probe and HST observations. *Icarus* 157, 373–400.
- Tody, D., 1993. Iraf in the nineties. In: Hanisch, R., Brissenden, R., Barnes, J. (Eds.), *Astronomical Data Analysis Software and Systems II*, A.S.P. Conference Ser., vol. 52, p. 173.
- Toigo, A., Gierasch, P.J., Smith, M.D., 1994. High resolution cloud feature tracking on Venus by Galileo. *Icarus* 109, 318–336.
- Tokumaru, P.T., Dimotakis, P.E., 1995. Image correlation velocimetry. *Exp. Fluids* 19, 1–15.
- Vasavada, A.R., Hörst, S.M., Kennedy, M.R., Ingersoll, A.P., Porco, C.C., Genio, A.D.D., West, R.A., 2006. Cassini imaging of Saturn: Southern-hemisphere winds and vortices. *J. Geophys. Res.*, 111.
- Vasavada, A.V., Ingersoll, A.P., Banfield, D., Bell, M., Gierasch, P.J., Belton, M.J.S., Klassen, K.P., Jong, E.D., Breneman, H., Jones, T.J., Kaufman, J.M., Magee, K.P., Senske, D.A., 1998. Galileo imaging of Jupiter's atmosphere: The Great Red Spot, equatorial region, and White Ovals. *Icarus* 135, 265–275.
- Yagi, G., Lorre, J., Jepsen, P., 1978. Dynamic feature analysis for Voyager at the image processing laboratory. In: *Proceedings of the Conference on Atmospheric Environment of Aerospace Systems and Applied Meteorology*, November 15–17, 1978. American Meteorological Soc., Boston, pp. 339–354.

<sup>3</sup> <http://archive.stsci.edu/hst/>.