doi:10.1068/p6820

Cast shadows in wide perspective

Sylvia C Pont, Maarten W A Wijntjes, Augustinus H J Oomes¶, Andrea van Doorn, Onno van Nierop¶, Huib de Ridder, Jan J Koenderink¶

Perceptual Intelligence Laboratory (π -lab), Industrial Design Engineering, Delft University of Technology, Landbergstraat 15, 2628 CE Delft, The Netherlands; e-mail: s.c.pont@tudelft.nl; ¶ Perceptual Intelligence Laboratory (π -lab), Electrical Engineering, Mathematics and Computer Science, Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands Received 10 September 2010, in revised form 16 July 2011

Abstract. We investigated the apparent spatial layout of cast shadows up to very wide fields of view. We presented up to 130° wide images in which two 'flat poles' were standing on a green lawn under a cloudless blue sky on a sunny day. The poles threw sharp cast shadows on the green, of which one was fixed. The observer's task was to adjust the azimuth of the shadow of the other pole such that it fitted the scene. The source elevation was kept constant. The two cast shadows are, of course, parallel in physical space, but generically not in the picture plane because of the wide perspective. We found that observers made huge systematic errors, indicating that, generically, they fail to account for these perspective effects. The systematic deviations could be well described by a weighted linear combination of the directions in the picture plane and in the physical space, with weights that depended on the positions of, and distance between, the poles.

1 Introduction

The geometry of shadows in perspective images has been a basic topic in arts and graphics education for centuries. So-called skiagraphy might be assumed to form common knowledge (eg Arnheim 1954; Baxandall 1995; Davies 1832; Descargues 1982; Gombrich 1995; Vicat Cole 1870). Generic cases of shadow projections might be quite intricate, eg the cases for complicated object shapes with arbitrary attitudes and heights that cast shadows on nonplanar surroundings. However, simple cases are easy to understand such as, for instance, cast shadows due to vertical grounded flat poles illuminated by a single collimated source (eg the sun), see figure 1: left for a source in front of the observer and right for a source behind the observer. A collimated source is a source whose rays are parallel. Cast shadows of vertical grounded objects due to a collimated source are parallel in 3-D space. This can be seen in views from above, see figures 2 and 3. However, in horizontal views these cast shadows converge towards the horizon because of perspective projection, similar to the convergence one may see in images of, for instance, railway tracks and streets extending in depth. The strength of the convergence or, in other words, the amount of perspective, is dependent on the distance of the viewer to the screen and on the distance of the screen to the objects in the scene (see figure 3).

Skiagraphy is common knowledge in computer graphics, where shadows are generally considered to be an important ingredient for the visualisation of spatial relationships between objects in a scene (Foley et al 1996). These relationships include the relative depth of objects in a scene and the height of objects with respect to a ground plane, and were acknowledged in early work on perception (Gibson 1950; Hochberg 1964). These ideas were confirmed formally in psychophysical experiments using computer renderings (Hubona et al 1999, 2004; Wanger et al 1992). In these studies observers had to adjust the sizes and positions of objects that were depicted on common screens. Projection and screen-size details were not mentioned by Hubona et al. Wanger et al used projection and viewing distances that coincided, with a stimulus size of 19.8 deg by 16.3 deg of visual angle.

Mamassian et al (1998) have extended such studies into the dynamic domain. They found that cast shadows were perceptually much more relevant for the recovery of spatial



Figure 1. Schematic depictions of perspective projections of cast shadows for flat poles on a flat ground plane under a single collimated source (eg the Sun): an example for a source (a) in front of the observer and (b) behind the observer. Note that the light rays from the Sun, and therefore the dashed helplines from the drawn Sun via the top corners of the poles to the top corners of the shadows, are parallel in physical 3-D space. The sides of the cast shadows are parallel in physical 3-D space as well. In these images they converge towards the horizon and the light rays seem to converge to a point, which effects are due to the perspective projection.



Figure 2. Photographs of cast shadows of trees with the Sun in front of the observer and at the back of the observer. The white dashed lines are drawn for visualisation of the shadow patterns. Note that in the tree pictures the shadows converge towards the horizon. The photograph of the camels was taken from above, at about a right angle to the ground surface, and shows that cast shadows of objects in sunlight are parallel in 3-D space. The convergence in the first two photographs is due to perspective.

arrangement and much less for object or background shape recovery. Allen (1999, 2000, 2001) showed that cast shadows provide important cues for distance judgments in real scenes and photographs of real scenes consisting of rods on or above a plane. He found that distance judgments became more veridical when cast shadows were present than when they were absent (Allen 1999). Using photocopied photographs in which images of cast shadows from rods and images of vertical rods were combined in inconsistent ways, he found results that suggest that observers are sensitive to cast shadow orientations if they have to judge distances (Allen 2000). However, from Allen's studies we cannot judge *how* sensitive human observers are to cast shadow orientations and also it was not clear what viewing angles he used.

On the other hand, studies on cast-shadow perception using visual search paradigms suggest that human observers are not very sensitive to cast-shadow direction. In these studies observers had to search for an odd-one-out in orthographically projected scenes of cylinders on a plane (Jacobson and Werner 2004) or of photographed pebbles (Lovell et al 2009). Porter et al (2010) also investigated cast-shadow perception via search paradigms and their findings suggest that observers find it hard to correlate objects casting shadows with their cast shadows. Porter et al stated that observers might seem insensitive to cast shadows because of difficulties distinguishing material from lighting effects (Kingdom 2008). Such object-shadow correspondence problems may be the main reason for the relative absence of cast shadows in Western art (Gombrich 1995) and can be used to speed up computations for Computer Graphics (Vangorp et al 2006).

In order to infer which shadows correspond with which objects, observers might use oversimplified strategies such as the copycat solution (Casati 2008). The copycat solution consists of a replica of the visible profile of the shadow-casting object, which in the norm yields an impossible shadow. If shadows in the outside (real shadows!) do not agree with such oversimplified processes they can even be regarded as 'anomalous'. For instance, mountain peaks always cast triangular shadows when the Sun is low and the observer is at the summit, *regardless* of the mountain profile. This effect was explained mathematically by Livingston and Lynch (1979). It is due to perspective, does not agree with oversimplified processes such as the copycat solution, and leads to effects that we cannot understand intuitively (or can even surprise us) and apparently need mathematical derivations to understand them. Why can perspective effects on shadow orientations be so non-intuitive? In this paper we address this question for a very simple case: a view of two vertical poles, with a spatial separation up to very wide viewing angles, on a flat horizontal surface, illuminated by a collimated source.

In figure 2 we show three examples of such simple cases in natural outdoors scenes. Informal observations made clear that even these images are hard to parse by many human observers. Note that these images were made under direct sunlight, the left photograph with the Sun in front of the observer, the middle with the Sun behind the observer and the right is a photograph taken from above, with the Sun behind the observer. Sunlight is almost perfectly collimated (the rays are nearly parallel) and thus sunlit vertical objects cast shadows on flat grounds that are parallel in the real world, which can be seen in the right image. However, the left and middle images show shadows that converge towards the horizon. This convergence is due to perspective and not to a diverging source. Human observers have difficulty in interpreting such converging (or diverging) directions in the picture plane as parallel in the scene. Such observations lead to questions about how human observers perceive cast-shadow orientations as a function of their separation in the visual field, and especially in cases of wide viewing angles.

Wide viewing angles occur under natural viewing conditions, in cinema (Cutting 1986), and currently commercially available equipment such as large television screens, projection systems, wide-angle lenses, etc. With the use of wide viewing angles, issues such as perspective transformations and the structure of visual space become important, even for viewing from the 'point of composition'. For instance, the majority of the population experiences an apparent visual field of only about 90°, much narrower than the dioptrics of the eye would suggest, namely a little over 180° (Koenderink et al 2009). Also, for a 130° viewing angle (in the same setup as is used in the current paper) the spatial attitudes of pictorial objects—cubes in general position—that look as having mutually parallel attitudes in pictorial space deviate up to 100° from the veridical attitudes (Koenderink et al 2010). Moreover, deviations were systematic: the cubes appeared parallel if they were similarly oriented with respect to the local visual rays and the local visual rays were subject-dependent.

In the current paper we address the question whether such systematic deviations from veridicality will also occur for cast-shadow orientations in a very simple scene under wide-angle viewing conditions? With 'simple' we mean that the scene is unambiguous in the sense that the ground plane and objects are flat, the cast shadows are explicitly declared to be cast shadows (not material/albedo variations) due to the flat poles (thus there is no shadow-object correspondence problem) and the single source

is given to be perfectly collimated (not another type of illumination). By 'wide angle' we mean a horizontal viewing angle of 130°, such that perspective effects are rather large. If cast shadows are an important ingredient for visualisation of spatial scene layout, it is expected that observers will be sensitive to their orientations. Despite such expected accuracy, we also expect systematic non-negligible deviations from the veridical shadow directions, as a function of position in the visual field (Koenderink et al 2009, 2010). We will test these hypotheses as a function of illumination orientation and as a function of the spatial separation of two shadow casting poles in a simple scene.

2 Experiment

We investigated the apparent spatial layout of cast shadows in wide perspective. Our experiment addresses the general question whether human observers account for the extent of the visual field when judging cast-shadow orientations on a flat ground plane.

2.1 Methods

2.1.1 *Participants*. Ten observers (six of them male, age ranging from 19 to 66 years), five naive and five of the seven authors participated in the experiment. The authors who participated in the experiment were not aware of the details of the stimulus conditions that were tested; the other two authors were and therefore did not participate as observers. All of them had normal or corrected-to-normal vision. The experiment was done in accordance with local ethical guidelines, Dutch Law, and with the Declaration of Helsinki.

2.1.2 *Stimuli*. The scene, see figure 3 for an example, was very simple and consisted of a green field with a blue sky above it, both 130° wide. On the green there were two vertical 'flat poles'. The scene was illuminated by a collimated beam, eg the Sun. Therefore the poles cast sharp shadows on the green.



Figure 3. Schematical representations of the experimental scene. (a) 3-D representation in which we show two poles on the green, the image projection plane (greyish, transparent), and the observer (sphere on a pole). Note that the shadows are parallel in 3-D space. (b) The corresponding image including dashed helplines from the cast shadows to the horizon and source, as in figure 1. The dashed lines are helplines for drawing the cast shadows in perspective projection under collimated illumination (the light rays are parallel in 3-D space but converge to a point in perspective projection). Note that, owing to the perspective projection, the shadows converge in this image.

The poles were placed at positions -60° , -30° , 0° , 30° , or 60° visual angle measured from the scene centre (the straight-ahead direction), with negative values denoting positions towards the left, see figure 4. Combining these positions results in 10 position combinations, which were all tested with the reference at one and test at the other position and vice-versa. The elevation of the source was kept constant at 45° (measured in 3-D space, with the 0° direction towards the viewer). The azimuth of the source could have 8 values, ranging from 22.5° to 337.5° , in steps of 45° . Altogether, this resulted



Figure 4. The conditions: poles could be positioned at -60° , -30° , 0° , 30° , or 60° visual angle measured from the centre (which is the straight-ahead direction), with negative values denoting positions towards the left.

in 160 trials, which were tested in random order in a single session, during which subjects could take a break as they wished. The images were rendered in real-time using custom-made software.

2.1.3 *Procedure.* The experiment was conducted on a Philips 82 cm wide flat display. The viewing distance was 19 cm and thus the scene subtended 130°. The observers viewed the stimulus monocularly, from a fixed vantage point directly in front of the scene's centre. We used a 'cheek-rest' for fixation (see figure 5). Because the observers need both lateral eye movements and minor head rotations about the vertical, special measures were taken to ascertain the constancy of the vantage point. The task was



Figure 5. The 'cheek-rest' with which we ascertained the vantage point.

explained with the following text in combination with the picture in figure 6:

"It is a crystal clear sunny day. You will see two 'flat poles' on a green lawn under a cloudless blue sky. Because it is so sunny and clear they throw sharp cast shadows. The shadow of one of the poles is fixed. Now adjust the shadow of the other pole such that it fits the scene. This doesn't fit!".



Figure 6. The task was explained to the observers by this image and the text in section 2.1.3.

2.1.4 *Results.* Figure 7 depicts the raw settings of the azimuthal angle as a function of the veridical azimuths. The 10 graphs represent the data for the ten observers. Each individual graph shows data for 160 trials. Points corresponding to a single condition



(see figure 4 for the 10 conditions) are connected separately for the reference left-test right and test left-reference right combinations. Thus, each plot contains 20 lines, each connecting 8 points. Overall, the data correlate quite well, with R^2 ranging from 0.86 to 0.96. Though these correlations are overall quite high, note that the individual settings show large errors: deviations of 80° are no exception. These visualisations also show that the deviations from veridical seem to be systematic. Therefore, we also correlated settings between observers, see figure 8a for a typical example. We found a range of 0.91 to 0.99 for R^2 . The R^2 s are, overall, higher for the between-observer correlations (grey bars in bargraph of figure 8a) than for the correlations between settings and veridical values (black bars). This finding confirms the suggestion that the errors that observers made are indeed systematic. A plot of the data in which the data points for the different conditions were connected (see figure 10 left graph for a typical example), qualitatively confirmed that the settings were dependent on the veridical azimuth in a non-linear way and on the condition.



Figure 8. (a) Scatterplot of a randomly chosen example of the settings of observer MT against the settings of observer HR. (b) The averaged R^2 s for each observer for the regressions against all other observers (grey bars), together with the R^2 s of the regressions of that observer's azimuthal settings against the veridical ones (black bars). These graphs clearly show that settings of observers correlate better with each other than with the veridical values.

2.1.5 Model evaluation. Detailed observations of images of observers' settings showed that the data were in between 'picture plane parallelism' and 'physical space parallelism', see figure 9. We show a representative graph of the raw data in figure 10a for subject HR. We fitted the data to a linear combination of 'picture plane parallelism' and 'physical space parallelism' via a weighted average. Let p denote the direction that is parallel in the picture plane and s the direction that is parallel in physical space. Let w denote a weight such that if w = 0 one has direction s, w = 1 direction p and



Figure 9. A schematic outline of how we fitted our data (d): we linearly combined 'picture plane parallelism' (p) and 'physical space parallelism' (s) via a weighted average.

w = 1/2 the bisectrix. Formally one has:

 $\boldsymbol{d} = w\boldsymbol{p} + (1 - w)\boldsymbol{s}.$

Figure 10b shows, for subject HR, for a single condition, the data (black points), the line corresponding to physical space parallelism (black line), picture plane parallelism (grey curves) and the fitted linear combination (dashed curves). The typical behaviour of the curves with 'veridical nodes' at 90° and 270° is due to the fact that these settings correspond with horizontal shadows in the picture plane.

Figure 11 shows the fitted weights as a function of condition, per observer (the 10 lines; the lines are for visualisation purposes only) and averaged over observers (the bars). We wanted to test whether our model fit describes the data sufficiently well and whether the use of different weights depending on conditions improves that description significantly. Therefore, we compared five candidate models via the Akaike Information



Figure 10. (a) The raw azimuthal settings connected per condition for observer HR, and (b) the corresponding fit for one condition. In 10a, the dashed lines represent the data that were modeled in the right graph (a single condition, with two lines for the reference left–test right and the test left–reference right data). In figure 10b, the grey curves depict picture plane parallelism, the black line physical space parallelism, and the dashed curves the fitted weighted average.



Figure 11. The model weights as a function of condition, per observer (the 10 lines; lines were drawn for the purpose of distinguishing data of individual observers). The bars depict the average weights as a function of condition. The condition numbers along the horizontal axis correspond to those with the depictions in figure 4.

Criterion (AIC) model selection using Akaike weights (Wagenmakers and Farrell 2004). The AIC is a method for comparing the adequacy of multiple, possibly non-nested models. The objective of AIC model selection is to estimate the information loss when the probability distribution associated with the true (generating) model is approximated by some probability distribution associated with the model under evaluation. The discrepancy between the true and approximating models is measured by the Kullback – Leibler information quantity. Instead of selecting a model on the basis of the difficult to interpret 'raw' AIC values, we analysed the data in terms of Akaike weights, which can easily be computed from the AIC values. The Akaike weights can be interpreted as the probability that a certain model is the best model (in the AIC sense that it minimises the Kullback – Leibler discrepancy), given the data and the set of candidate models (Wagenmakers and Farrell 2004).

The five candidate models/hypotheses were that observers adjust the shadow azimuths: (1) parallel in physical space (*S*, veridical);

- (2) parallel in the picture plane (*P*);
- (3) as a weighted average over S and P (1 free parameter);
- (4) as a weighted average over S and P with the weight depending on the conditions that were grouped in the 5 rows in figure 4 (5 free parameters);
- (5) as a weighted average over S and P with the weight depending on the conditions in figure 4 (10 free parameters).

For each observer we got a set of 5 Akaike weights, which sum up to 1 (by definition). We found Akaike weights which were negligible (smaller than 10^{-3}) for models 1 and 2, between negligible and 0.11 for model 3, between 0.83 and 0.99 for model 4, and between 0.01 and 0.17 for model 5. Thus, the fourth model explained the data best (of these 5 models) for *all* ten observers. This model describes the data as a weighted average over physical space parallelism and picture plane parallelism, which depends on the absolute positions of the poles.

3 General discussion

Our stimulus consists of a simple scene, illuminated by a single collimated source, with a flat ground plane on which there are two flat poles casting shadows on the ground plane. One of the cast shadows is given, and observers are able to adjust the shadow of the other pole such that it fits the scene. The settings correlate quite well with the veridical values (0.86 to 0.96). However, the inter-observer correlations are even higher (0.91 to 0.99). Moreover, the deviations show a systematic pattern, as could be seen in figures 10 and 11. The deviations depend on the absolute pole positions and on the difference between the veridical direction and parallel direction in the picture plane. Errors up to 80° were found to be quite common in this setting.

We did not find different results for the authors and for the naive observers. All observers showed the same qualitative pattern. It would be interesting to test very experienced observers though, such as drawing teachers, painters, and graphics artists. Furthermore, the scene in this experiment was very simple and it might be interesting to see whether the addition of cues for the geometry of the scene would decrease the deviations from veridical. One could think of texture on the ground plane, additional objects in the scene, adding shading and reflexes, using realistic scenes, etc. In object positioning and sizing experiments, such cues had, overall, no large effects on performance (Hubona et al 1999, 2004; Wanger et al 1992). We expect that the effects of such cues on our outcomes will also be negligible (eg see figure 2).

The data agree well with a linear weighted average of the directions for picture plane parallelism and physical space parallelism. Most observers adjusted the cast shadow in between these two cases such that the average weights were found to be in the range 0.27 to 0.67. The weights depend significantly on the positions of and

distances between the poles. This suggests that the converging cast shadows in figure 2 are judged to be converging instead of parallel shadows. These results are in agreement with our previous findings of effects of wide-angle views (Koenderink et al 2009, 2010). Our results seem to stem from a misjudgment of the extent of the scene. In other words, most human observers to some degree fail to account for the divergence of visual directions.

It remains to be investigated how this finding might be related to light-field perception. The light field, essentially the radiance, captures the luminous environment or everything that we can potentially see. Human observers are sensitive to the physical light field in the sense that they are able to adjust the direction, diffuseness, and intensity of a matte white sphere (a so-called gauge object) introduced at arbitrary locations in photographs of a large variety of scenes (Koenderink et al 2007; Pont et al 2010). That we were able to measure the 'visual light field' was no guarantee that the visual light field will be consistent in the radiometric sense that the physical light field is. In the present study we have found that adjusted cast shadow orientations due to collimated light within a single light zone (Madsen 2006) or framework of illumination (Gilchrist and Radonjic 2010) deviate from the veridical orientations. These deviations vary systematically as a function of global illumination direction and of pole positions in the visual field. Thus, estimated cast shadow orientations are not consistent with the global direction of illumination (or parallelism in physical space). The estimations deviate toward parallelism in picture space and do not have a vantage point-a striking effect that appears to be undocumented in the literature.

Acknowledgments. Maarten Wijntjes was supported by the Netherlands Organisation for Scientific Research (NWO).

References

- Allen B P, 1999 "Shadows as sources of cues for distance of shadow-casting objects" *Perceptual* and Motor Skills **89** 571-584
- Allen B P, 2000 "Angles of shadows as cues for judging the distance of shadow casting objects" Perceptual and Motor Skills **90** 864–866
- Allen B P, 2001 "Lighting position and judgments of distance of shadow-casting objects" *Perceptual* and Motor Skills **93** 127-130
- Arnheim R, 1954 Art and Visual Perception (Berkeley, CA: University of California Press)
- Baxandall M, 1995 Shadows and Enlightenment (New Haven, CT & London: Yale University Press)
- Casati R, 2008 "The copycat solution to the shadow correspondence problem" Perception 37 495-503
- Cutting J E, 1986 "The shape and psychophysics of cinematic space" *Behavior Research Methods, Instruments, & Computers* 18 551-558
- Davies C, 1832 A Treatise on Shades and Shadows and Linear Perspective (New York: J and J Harper)
- Descargues P, 1982 Perspective-History, Evolution, Techniques (New York: Van Nostrand Reinhold)
- Foley J D, van Dam A, Feiner S K, Hughes J F, 1996 Computer Graphics—Principles and Practice (New York: Addison-Wesley)
- Gibson J J, 1950 The Perception of the Visual World (Boston, MA: Houghton Mifflin)
- Gilchrist A L, Radonjíc A, 2010 "Functional frameworks of illumination revealed by probe disk technique" *Journal of Vision* **10**(5):6, 1–12
- Gombrich E H, 1995 The Depiction of Cast Shadows in Western Art (London: National Gallery)
- Hochberg J E, 1964 Perception (London: Prentice-Hall International)
- Hubona G S, Wheeler P N, Shirah G W, Brandt M, 1999 "The relative contributions of stereo, lighting, and background scenes in promoting 3D depth visualization" ACM Transactions on Computer-Human Interaction 6 214–242
- Hubona G S, Shirah G W, Jennings D K, 2004 "The effects of cast shadows and stereopsis on performing computer-generated spatial tasks" *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans* 34 483–493
- Jacobson J, Werner S, 2004 "Why cast shadows are expendable: Insensitivity of human observers and the inherent ambiguity of cast shadows in pictorial art" *Perception* **33** 1369–1383
- Kingdom F A A, 2008 "Perceiving light versus material" Vision Research 48 2090-2105
- Koenderink J J, Pont S C, van Doorn A J, Kappers A M L, Todd J T, 2007 "The visual light field" *Perception* 36 1595-1610

- Koenderink J J, Doorn A J van, Todd J T, 2009 "Wide distribution of external local sign in the normal population" *Psychological Research* 73 14–22
- Koenderink J J, van Doorn A J, de Ridder H, Oomes A H J, 2010 "Visual rays are parallel" Perception 39 1163-1171
- Livingston W, Lynch D, 1979 "Mountain shadow phenomena" Applied Optics 18 265-269
- Lovell P G, Gilchrist I D, Tolhurst D J, Troscianko T, 2009 "Search for gross illumination discrepancies in images of natural objects" *Journal of Vision* **9**(1):37, 1–14
- Madsen M, 2006 "Light-zones: as concept and tool; An architectural approach to the assessment of spatial and form-giving characteristics of daylight" *Proceedings International Conference on Architectural Research ARCC/EAAE* (Philadelphia, PA: ARCC)
- Mamassian P, Knill D C, Kersten D, 1998 "The perception of cast shadows" Trends in Cognitive Sciences 2 228-295
- Ostrovsky Y, Cavanagh P, Sinha P, 2005 "Perceiving illumination inconsistencies in scenes" Perception 34 1301-1314
- Pont S C, van Doorn A J, de Ridder H, Koenderink J J, 2010 "The visual light field in the in and outside" *Perception* **39** Supplement, 104
- Porter G, Tales A, Leonards U, 2010 "What makes cast shadows hard to see?" *Journal of Vision* **10**(3):13, 1–18
- Vangorp P, Dumont O, Lenaerts T, Dutré P, 2006 "A perceptual heuristic for shadow computation in photo-realistic images", in SIGGRAPH'06: ACM SIGGRAPH 2006 Sketches 102 (New York: ACM) http://dl.acm.org/citation.cfm?doid=1179849.1179977
- Vicat Cole R, 1870 Perspective as Applied to Pictures (London: Seeley, Service & Co)
- Wagenmakers E-J, Farrell S, 2004 "AIC model selection using Akaike weights" Psychonomic Bulletin & Review 11 192-196
- Wanger L R, Ferwerda J A, Greenberg D P, 1992 "Perceiving spatial relationships in computergenerated images" IEEE Computer Graphics & Applications 12(3) 44–58

ISSN 1468-4233 (electronic)



www.perceptionweb.com

Conditions of use. This article may be downloaded from the Perception website for personal research by members of subscribing organisations. Authors are entitled to distribute their own article (in printed form or by e-mail) to up to 50 people. This PDF may not be placed on any website (or other online distribution system) without permission of the publisher.