A Virtual Try-On System for Prescription Eyeglasses

Qian Zhang and Yu Guo ■ *Nanyang Technological University*

Pierre-Yves Laffont, Tobias Martin, and Markus Gross ■ *ETH Zurich*

Corrective lenses introduce distortion caused by the refraction effect, which changes the wearer's appearance. To give users a more realistic experience, a virtual try-on system modifies an input video and virtually inserts prescription eyeglasses with reflections and shading, producing an output similar to a virtual mirror.

Vision-correcting eyeglasses have improved the lives of millions of people. Eyeglasses significantly affect the wearer's appearance, and the selection of new pairs of eyeglasses is largely based on how the glasses look when wearers try them on. However, an often overlooked fact is that corrective lenses introduce distortion caused by the refraction effect. As Figure 1 illustrates, the eyes of a person wearing corrective lenses for nearsightedness appear smaller compared with wearing nonprescription lenses, whereas the eyes of a person wearing lenses for farsightedness appear larger.

The traditional process of trying on and picking new eyeglasses frames in a brick-and-mortar shop has a significant shortcoming: eyeglasses on the display are equipped with demo lenses that have zero corrective power, and thus refraction does not deform the eyes. Thus, customers cannot see what they will actually look like until their custom prescription lenses are installed in the frames and the sale is final. Their appearance will differ from the in-store trial, which may cause disappointment and buyer's remorse, especially for customers with strong eyeglasses prescriptions. A similar issue occurs with online stores, which allow customer to virtually try-on eyeglasses frames by overlaying them onto an input image. The online systems still do not adjust the image for the refraction effect.

We present a system for virtually trying on prescription eyeglasses. Our system acts as a virtual mirror, allowing users to try on a variety of eyeglasses with corrective lenses based on their prescription (see Figure 2). We use an image sequence of the user without eyeglasses as input, along with the user's eyeglasses prescription and a 3D model of the desired eyeglasses frame. Our system generates a 3D representation of the corrective lenses mounted into the eyeglasses frame and modifies the video sequence to virtually insert the eyeglasses using image-based rendering. This approach simulates the distortion introduced by the prescription lenses and gives users a better idea of how they would look when wearing the new pair of eyeglasses.

To the best of our knowledge, the proposed virtual try-on system for prescription eyeglasses is the first to account for refraction effects. (See the "Related Work in Virtual Try-On Applications" sidebar for more details.) Our system was inspired by the traditional eyeglasses manufacturing pipeline followed by opticians. We generate a 3D representation of the corrective lenses that fit the user's eyeglasses prescription and the chosen eyeglasses frame. Then, an image-based rendering technique virtually inserts prescription eyeglasses into the input video, while taking into account the effects of refraction, reflection, and shading. The findings from our user study highlight the importance of refraction and reflection in the perceived realism of virtual try-on results.

System Overview

The virtual try-on system we developed inserts prescription eyeglasses onto the user's face and

 Published by the IEEE Computer Society

simulates important changes to the appearance due to refraction, reflection, or shadows cast on the face. The system uses the following three elements as input:

- *Image sequence*. An image sequence of the user without eyeglasses is captured with a color camera.
- *User's eyeglasses prescription*. An eyeglasses prescription, usually provided by an optometrist, specifies the value of all parameters necessary to correct blurred vision due to refractive errors, including myopia, hyperopia, presbyopia, and astigmatism. Table 1 shows a typical eyeglasses prescription.
- *Eyeglasses frame*. The user chooses the desired eyeglasses frame. The eyeglasses geometry is typically accessible from online stores, which scan and digitize the eyeglasses frames. For this work, we purchased 3D models for six different commercially available eyeglasses frames from TurboSquid (www.turbosquid.com/3d-model/glasses/).

Figure 3 gives an overview of our approach's pipeline, which consists of two main stages: virtual eyeglasses generation and video synthesis.

In the first stage, we generate a 3D representation of the prescription eyeglasses (the frame and corrective lenses), with an appropriate position relative to the user's face geometry. Inspired by the traditional eyeglasses manufacturing pipeline, this virtual eyeglasses generation stage has three steps:

1. *Positioning the eyeglasses on the user's face*. After an initial manual positioning step for the first frame, we use face tracking to automatically align the eyeglasses with the user's face in the following frames.
2. *Creating a parametric lens model.* Based on the user's prescription and desired lens properties, this model describes the geometry of the uncut lens before mounting.
3. *Cutting and mounting the lens.* We trim the lens geometry according to the shape of the eyeglasses frame and insert the virtual lenses into the eyeglasses frame.

In the video synthesis stage, we render the virtual eyeglasses and insert them into the input image sequences, taking into account the eyeglasses frame, lenses, and surrounding lighting. In doing so, we account for the effects of refraction, reflection, and shadows due to the inserted eyeglasses.

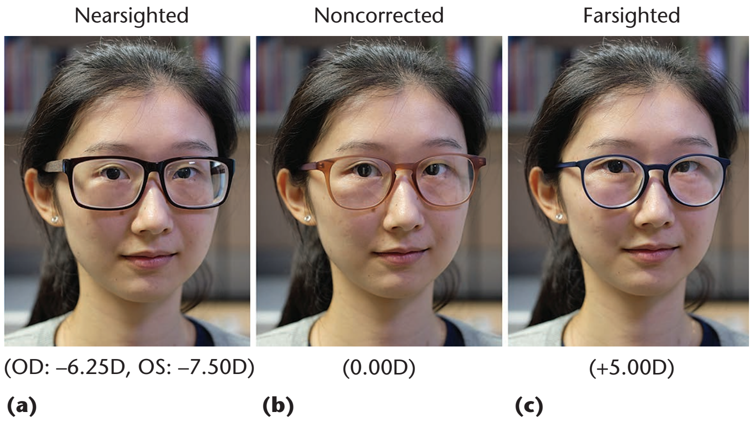

Figure 1. Prescription eyeglasses introduce refraction effects that change the wearer's appearance. The eyes of a person wearing corrective lenses for nearsightedness (a) will appear smaller compared with (b) wearing nonprescription lenses, whereas the eyes of a person wearing lenses for farsightedness (c) will appear larger. The numbers below each image correspond to the optical power of the lenses measured in diopters, where OD and OS denote the right and left eye, respectively.

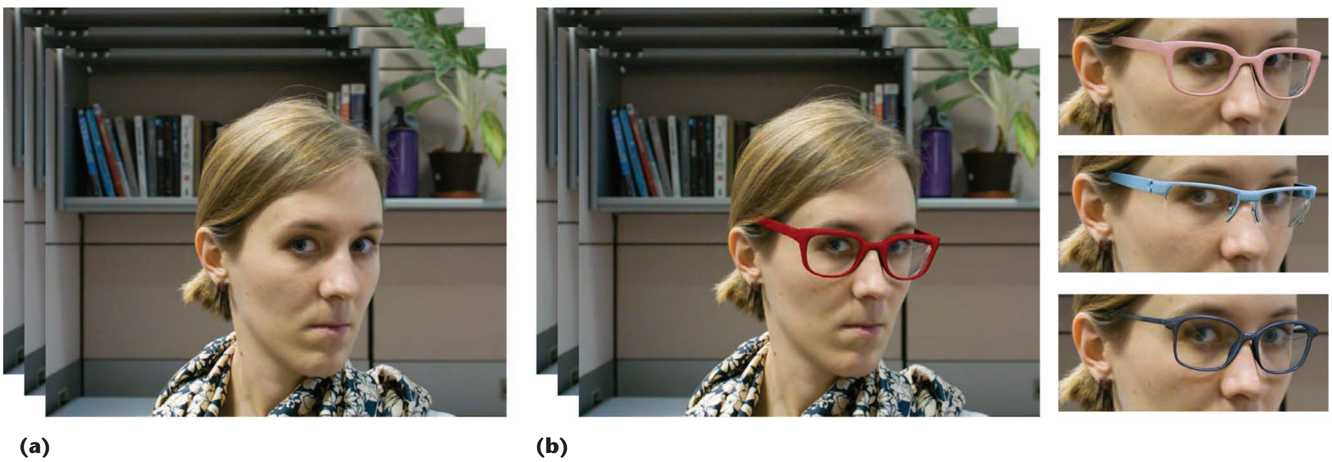

Figure 2. Our virtual try-on system for prescription eyeglasses modifies (a) an input video and (b) virtually inserts prescription eyeglasses, producing an output similar to a virtual mirror. Our approach handles reflection and shadows and simulates the effects of refraction based on the user's eyeglasses prescription. A variety of eyeglasses frames are available for selection.

Related Work in Virtual Try-On Applications

Augmented-reality technologies allow us to integrate virtual objects into real-world video sequences, enabling computer-generated objects to be inserted into an input image or video as if they were part of the observed scene.[1] These technologies enhance people's perception of reality and enable a variety of applications in fields such as education, maintenance, design, and e-shopping. Virtual try-on systems make the previsualization of products possible and provide users with a more credible try-on experience from the comfort of their own homes. Researchers have proposed virtual try-on systems for objects such as clothing[2–4] and eyeglasses.[5,6]

An eyeglasses virtual try-on application inserts virtual eyewear, such as vision-correcting eyeglasses frames or sunglasses, onto a user's face, which is captured by a color camera. Many online eyeglasses stores allow users to upload a frontal face image and insert glasses on it, including GlassesUSA (www.glassesusa.com) and JCPenney (www.jcpenneyoptical.com/virtual-try-on/). Other commercial applications insert eyewear directly into a live video stream captured by webcam; examples include FittingBox (demo.fittingbox.com/fitlive/single-example.html) and SmartBuyGlasses (www.smartbuyglasses.com/virtual-try-on). Research literature on virtual eyeglasses try-on proposes image-based blending techniques to insert eyeglasses onto the user's face.[7,8] Those existing systems, either using still images or video, act as a virtual mirror that lets users try on eyeglasses virtually. However, to the best of our knowledge, all available virtual try-on solutions ignore the refraction effects caused by eyeglasses lenses.

In this article, we demonstrate that refraction artifacts, which occur with real prescription lenses, drastically change the wearer's appearance, thus affecting how customers perceive their potential new pair of eyeglasses. The proposed system acts as a virtual mirror, but unlike previous methods, it leverages 3D geometry to simulate refraction, reflection, and shadows, yielding virtual try-on results with increased perceived realism.

References

1. D. Van Krevelen and R. Poelman, "A Survey of Augmented Reality Technologies, Applications and Limitations," *Int'l J. Virtual Reality*, vol. 9, no. 2, 2010, pp. 1–20.
2. S. Giovanni et al., "Virtual Try-On Using Kinect and HD Camera," *Proc. Int'l Conf. Motion in Games*, 2012, pp. 55–65.
3. N. Magnenat-Thalmann et al., "3D Web-Based Virtual Try On of Physically Simulated Clothes," *Computer-Aided Design and Applications*, vol. 8, no. 2, 2011, pp. 163–174.
4. M. Cao et al., "Educational Virtual-Wear Trial: More Than a Virtual Try-On Experience," *IEEE Computer Graphics and Applications*, vol. 35, no. 6, 2015, pp. 83–89.
5. D. Tang et al., "Making 3D Eyeglasses Try-On Practical," *Proc. IEEE Int'l Conf. Multimedia and Expo Workshops* (ICMEW), 2014; doi:10.1109/ICMEW.2014.6890545.
6. S.-H. Huang, Y.-I. Yang, and C.-H. Chu, "Human-Centric Design Personalization of 3D Glasses Frame in Markerless Augmented Reality," *Advanced Eng. Informatics*, vol. 26, no. 1, 2012, pp. 35–45.
7. O. Déniz et al., "Computer Vision Based Eyewear Selector," *J. Zhejiang Univ. SCIENCE C*, vol. 11, no. 2, 2010, pp. 79–91.
8. M. Afifi and M. Korashy, "Eyeglasses Shop: Eyeglasses Replacement System Using Frontal Face Image," *Proc. 4th Int'l Conf. Mathematics and Information Science* (ICMIS), 2015.

Table 1. Example eyeglasses prescription.*

Eye	Sphere	Cylinder	Axis	PD
OD	–4.25	–0.75	160	64
OC	–4.50	–	–	

**OD and OS represent the lens prescription of the right and left eye (from the wearer's point of view), respectively. Sphere and cylinder are the spherical and cylindrical correction, and axis indicates the cylinder axis in the case of astigmatism-correcting lenses. Pupillary distance (PD) is the distance between the pupil centers.*

Virtual Eyeglasses Generation

To begin our discussion, we briefly describe the traditional eyeglasses manufacturing process, before introducing our proposed system.

Once the customer has chosen an eyeglasses frame for purchase, the optician measures the pupillary distance (PD), which is the horizontal distance between the left and right pupils. This can be done by marking the position of the pupils on the demo lenses, while the customer has the glasses on. This step is essential to ensuring that the prescription lenses will be appropriately positioned with respect to the eyes.

The next step is to choose lens blanks based on the strength of the correction needed and desired lens properties (for example, lens material). Lens blanks are circular, uncut lenses that are usually stocked by the lens manufacturers, with a variety of front surface curvatures. If necessary, the back surface of the lens is ground and polished to produce a lens according to the desired prescription.

The eyeglasses frame is then inserted into a dedicated tracing machine in order to measure its inner contours, which will be used to cut the lens blanks to the appropriate shapes. Each lens blank is placed into an instrument to locate and mark its optical center; these points will be positioned in front of the customer's pupils to ensure optimal vision. Finally, an edging machine is used to trim the lens blanks into the proper lens shapes, according to the previously measured contours. The cut lenses are then inserted

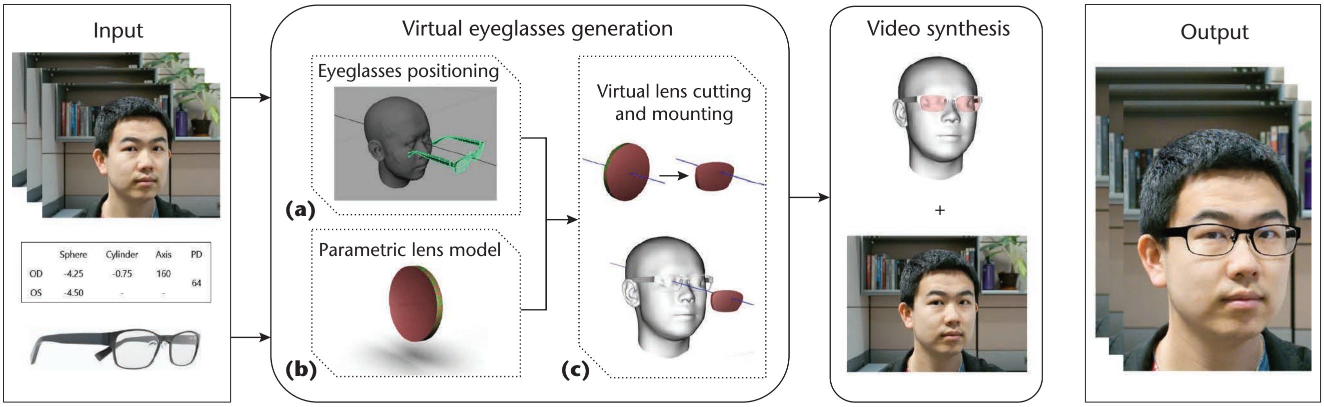

Figure 3. Overview of the pipeline for virtual try-on of prescription eyeglasses. Our system takes as input a video of the user, the user's eyeglasses prescription, and a 3D model of the desired eyeglasses frame. In the virtual eyeglasses generation stage, we first create a virtual 3D representation of the desired prescription eyeglasses, by (a) positioning the eyeglasses frame with respect to the user's face, (b) building a parametric lens model based on the user's eyeglasses prescription, and (c) cutting the corrective lenses before mounting them into the eyeglasses frame. In the second stage, video synthesis, we use image-based rendering to generate a synthetic image sequence where the prescription eyeglasses are virtually inserted, taking into account the effects of refraction, reflection, and shadows due to the inserted eyeglasses.

into the eyeglasses frame.

We create virtual eyeglasses with a similar process. First, we place the eyeglasses frame appropriately onto the user's face geometry. Next, we build a parametric model representing the geometry of each lens according to the user's eyeglasses prescription. Finally, lenses are cut and mounted into the eyeglasses frame.

Eyeglasses Positioning

Similar to the optician pipeline, we first place the eyeglasses frame with respect to the user's face geometry. We obtain the geometry and pose of the user's face for each frame by tracking the face using the Faceshift software[1] and a Primesense Carmine 1.09 RGBD sensor. Calibration between the RGBD sensor and the color camera, which is used to capture the user's input image sequence, is performed with a camera calibration toolbox.[2] The camera's intrinsic and extrinsic parameters let us align the face geometry with the input color images.

Next, we manually position the eyeglasses onto the face mesh for the first frame. For all the examples we tested, this process took less than 5 minutes on average. A fully automatic option would use affine transformation computed based on preselected feature points on face and eyeglasses 3D model[3] or a physics-driven technique.[4]

After the initial manual positioning of the eyeglasses for the first frame, we track the head pose to automatically align the eyeglasses with the user's face in the subsequent frames. This is achieved by calculating the relative pose change in each frame.

Parametric Lens Model

Given the user's eyeglasses prescription, we generate the 3D lens geometry based on a parametric model so that the optical power of the virtual lens corresponds to the user's prescription. A lens is a 3D transparent and closed object. It consists of two main surfaces: a front surface and a back surface. The lens thickness is defined as the distance between front and back surface along its optical axis. Physical lenses are made of a transparent material with a certain refraction index, which affects lens thickness, weight, and optical properties.

Optical power refers to a lens' ability to bend light rays, as specified by the eyeglasses prescription. The front and back surface curves determine the lens' optical power. Spherical lenses are rotationally symmetric, and their front and back surfaces have a constant curvature. In contrast, the surface curvature of toroidal lenses, which are used to correct astigmatism, varies with the direction; it is usually defined along two orthogonal directions called the axis meridian and power meridian. Modern lenses generally take a meniscus shape, with convex front curves and concave back curves. The optical power P of a lens measured in diopters is given by

$$P = F + B + (t/h) * F^2, \tag{1}$$

where F and B are the front and back power in diopters, t is the lens center thickness in meters, and η is the index of refraction. The focal power P is specified by the user in the form of an eyeglasses prescription (see Table 1).

Multiple lenses can achieve the same correction, so the user can choose among different materials (with a different refractive index η), thicknesses t (which depends on the material), and prices. Similar to the manufacturing process, we choose an appropriate base curve for the front lens surface based on the optical power P of the lens. The lens's base curve is the surface curve that becomes the basis from which the remaining curves will be calculated. For modern ophthalmic lenses, the base curve is typically the front curve of the lens blank, which has a front power F. Base curve selection charts,[5] which are available from manufacturers, provide the recommended ranges of final surfaced power for each base curve in the series. Knowing the optical power P, front power F, lens center thickness t, and index of refraction η, we can calculate the back surface power B using Equation 1. Once all the quantities are known, we can generate the 3D geometry of the lens.

For the sake of simplicity, we only describe the lens model for spherical lenses here. However, we can also generate lenses for astigmatism and presbyopia as well as bifocal and progressive lenses, given the user's prescription.

Virtual Lens Cutting and Mounting

Inspired by real lens-cutting machines, we detect the 2D inner contour of the eyeglasses frame from a front-facing view and extrude that contour to cut the lens. In the process, the uncut lens is aligned with the optical axis of the eye, making sure that the lens optical center sits in front of the pupil. The cut lens is represented using a triangle mesh with a fine tessellation. After the lens cutting, we insert each corrective lens into the eyeglasses frame by translating it along its optical axis.

Video Synthesis

In the next stage, we insert the virtual eyeglasses into the input sequence using image-based rendering, where the eyeglasses are rendered using ray tracing. From the previous stage, we obtained a parametric lens model, the well-positioned eyeglasses, and the user's face geometry for each image frame. We address the rendering process by first describing the objects in the virtual scene and the materials associated with them. Then, we describe the ray-tracing rendering of the lenses, with the refraction and reflection effects and the shading cast on the user's face.

Scene Description

We first prepare a virtual scene for the rendering process, where the user is wearing the prescription eyeglasses. Obtained from the previous steps, there are four objects in the scene: two corrective lenses, the eyeglasses frame, the user's face mesh, and the background. In ray tracing, each primary ray traced from the camera and intersecting the scene geometry is assigned a color, which depends on the local material and shading (that is, the quantity of light received at the intersection point). Primary rays that do not intersect any of the scene objects are assigned the same color as in the input image.

To produce plausible shading, we first build a representation of the incident lighting that surrounds the scene. A simple way to capture this is to place a chrome sphere in the scene, before capturing the input video. A photograph of the sphere is then unwrapped into an environment map,[6] which represents the amount of incident light coming from every direction. Other approaches for estimating the incident lighting include automatic methods such as shape-from-shading[7–9] and the real-time method.[10] In this system, we utilize the chrome sphere. Alternatively, we could use a precaptured default environment map.

Each object in the scene has an associated material. For virtual objects, the material is set by default or preselected by users. The color on the eyeglasses frame is determined using Phong shading,[11] the properties (specularity) of which can be adjusted to change the eyeglasses color. The lenses are associated with a dielectric material that refracts or reflects light rays. For the user's face mesh and background, the materials come from the input image sequences. The user's face is considered a diffuse surface. To insert plausible shadows while still preserving the details from the input image sequence, we employ an image-based approach.

Ray Tracing

We render the scene using a ray-tracing-based technique, which takes into account the refraction and reflection effects introduced by the corrective lenses and the shadow cast on the face by the eyeglasses frames. Rays emanating from the camera position are cast through pixel centers of the images and intersected with the scene (Turner-Whitted style[12]).

Corrective lenses introduce refraction effects and distort regions behind the lens. In our system, lenses are rendered using dielectric materials, which consist of a refraction part and an internal reflection part.[13] Rays are refracted using Snell-Descartes' law, which is a function of ray direction, the surface normal at the intersection point, and the lens refraction index. As it enters the first lens surface, each ray is refracted or reflected, either into the lens or back into the scene (see Figure 4). Total internal reflection occurs for some of the rays entering the lens, where they keep bouncing back and forth within the lens.

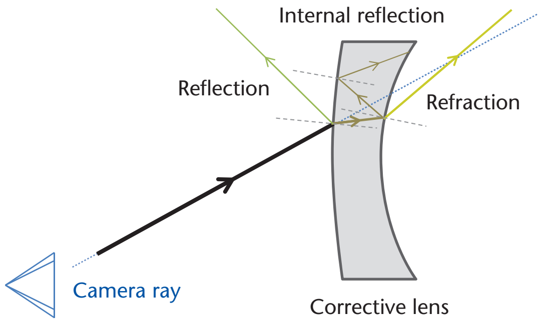

Figure 4. Ray tracing for simulating the refraction and reflection effects. As camera rays (blue) enter the first surface of a corrective lens, each ray is refracted or reflected, either into the lens or back into the scene. For some of the rays entering the lens, total internal reflection occurs.

We simulate reflection effects on the surface of the transparent lenses by leveraging the environment map. When a ray is reflected by the lens surface and does not further intersect any scene object, an environment map texture lookup is performed based on the ray's direction.

To simulate the shadows cast by eyeglasses, we estimate the shading at each visible point on the face using Monte Carlo integration.[14] The algorithm in Figure 5 describes this procedure.

For each point p at which we wish to estimate the shading, an integration over a hemisphere is performed. The hemisphere is centered on the local surface normal of point p and sampled using importance sampling[14] with weight w_i. We cast rays originating from p toward sample points on the hemisphere (see Figure 6). Each ray is then tested for intersections with objects in the scene. Rays that do not intersect with any object contribute to the local shading (we can look up their color in the environment map), whereas rays that are blocked by occluders do not. We perform the integration twice: once in a virtual scene containing solely the face geometry, yielding shading $S_{noGlasses}(p)$, and once with the added eyeglasses frame, yielding shading $S_{withGlasses}(p)$. We obtain the final color $I_{withGlasses}$ by multiplying the input color $I_{noGlasses}$ with the shading ratio:

$$I_{withGlasses}(p) = (S_{withGlasses}/S_{noGlasses}) * I_{noGlasses} \quad (2)$$

This process ignores lens occlusion and lens effects, such as caustics, which significantly speeds up rendering time.

Lastly, to account for users with long hair and inaccuracies in the estimated face geometry, we smoothly fade out the eyeglasses frames near the ear region. We blend each pixel of the output image with the input using image compositing, based on the distance to the front of the face.

Input: Intersection points p between camera rays and face geometry; color $I_{noGlasses}(p)$ from RGB image.
for each point p **do**
  $S_{noGlasses} = (0, 0, 0)$
  $S_{withGlasses} = (0, 0, 0)$
  Sample the hemisphere at p with weight w_i
  for each sample point q_i in the hemisphere **do**
    ray direction $\vec{d}_i = \left(\overrightarrow{pq_i}\right) / \left\|\overrightarrow{pq_i}\right\|$
    $S_i = w_i * \mathrm{EnvMap}\left(\vec{d}_i\right)$
    $S_{noGlasses} \mathrel{+}= S_i$
    if ray $\overrightarrow{pq_i}$ does not hit eyeglasses frame **then**
      $S_{withGlasses} \mathrel{+}= S_i$
    end if
  end for
  $I_{withGlasses}(p) = (S_{withGlasses}/S_{noGlasses}) * I_{noGlasses}$
end for
Output: color $I_{withGlasses}(p)$ at each intersection p.

Figure 5. Shading estimation algorithm. We estimate the shading at each visible point on the face using Monte Carlo integration.

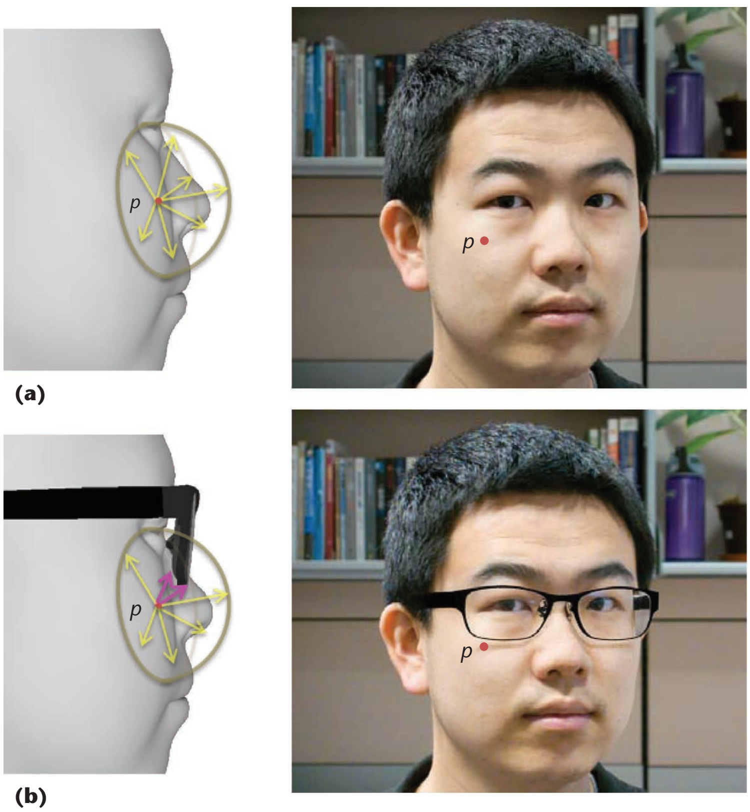

Figure 6. Monte Carlo integration process for shading computation. (a) To estimate the shading at a point p, rays are traced from p to several directions of the viewing hemisphere. (b) Once virtual eyeglasses are inserted, the eyeglasses frame blocks some of these rays (rays in magenta). This results in shadows cast on the user's face, which our approach is able to simulate.

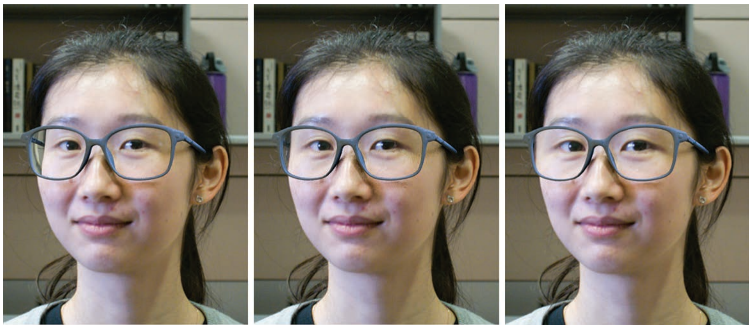

Figure 7. Comparison of our synthetic results, where eyeglasses are mounted with different lenses: (a) nearsighted, (b) noncorrected, and (c) farsighted. Similar to the photographs in Figure 1, the eyes appear smaller behind nearsightedness lenses, and they are magnified by farsightedness lenses. This phenomenon becomes more obvious as the strength of the correction increases.

Results

The distortion introduced by prescription eyeglasses becomes more obvious as the strength of the prescription increases. Similar to the changes that occur when people wear real eyeglasses (see Figure 1), our synthesized images alter the size of the eyes (see Figure 7). Eyes appear smaller behind nearsightedness lenses, and they are magnified by farsightedness lenses. In addition to the eyeglasses frames themselves, this important phenomenon changes how people look when they put on glasses.

Figure 8 compares eyeglasses virtual try-on results. Given the input image (Figure 8a), the virtual try-on solution from an online eyeglasses store (Figure 8b) only inserts the eyeglasses frame, without considering the lenses. In contrast, our virtual try-on result (Figure 8c) takes into account refraction, reflection, and shadows, so the results are more similar to the real reference image (Figure 8d) captured with a similar pose and eyeglasses for the same prescription. Our simulation of the distortion due to corrective lens refraction gives the user a more realistic experience.

The supplemental online video (see https://youtu.be/_fckwZCzCgc) provides results for several other virtual try-on sequences, including multiple users and eyeglasses frames.

User Study

We performed a user study to assess the perceived realism of virtual try-on videos generated with our approach. Participants in the user study were shown multiple videos of virtual try-on results and asked to rank them based on their perceived realism.

To generate the stimuli, we captured videos using five actors of different ethnicities and one mannequin. In each capture session, we asked the actor to try on prescription eyeglasses that we provided, with sphere powers ranging from –1 to –7.5 diopters. We recorded this as the reference video. We then recorded a second video of the actor without eyeglasses and used this video as input to our virtual try-on. We generated four variations of virtual try-on results using our approach, with/without reflections and with/without refraction (Figure 9), using lens model parameters corresponding to the prescription lenses in the reference video. We also recorded a live virtual try-on session using the SmartBuyGlasses virtual try-on online eyeglasses store (www.smartbuyglasses.com/virtual-try-on) with similar eyeglasses frames. The five stimuli videos were cropped to less than 3 seconds for each actor.

In each trial of the study, we first showed the participants the reference video of an actor wearing real eyeglasses. The five stimuli corresponding to this actor were then shown simultaneously, in a loop. We asked the participants to rank the videos according to how they looked, from most to least realistic, by dragging them over the screen into ranking bins. Each trial corresponding to a single actor was repeated twice, and all trials were randomly ordered. Two training trials were added at the beginning of the session but were not used in the analysis. In total, each participant completed 14 trials.

Twenty individuals (12 male and eight female) participated in our study, with ages ranging from 23 to 49 (average 28). Nine participants reported that they were familiar with computer graphics, and 17 of them wore eyeglasses. The average length of the study (including a short break) was 16 minutes.

Based on the ranking results, we assigned scores to all the videos. In each trial, the best-ranked video was assigned a score of 5, and the lowest-rank video was assigned a score of 1. Upon analyzing the scores from all trials and all participants, the results show that the participants overall favored video sequences that exhibit refraction and reflections, which increased the perceived realism in our virtual try-on results (see Figure 10 and Table 2).

Table 2 shows that 62.50 percent of the votes favored our videos with refraction compared with our videos with no refraction, for sequences with reflections. A paired t-test confirmed a statistically significant difference ($t = 7.09$, $p < 0.005$). Also, a similar analysis can be made when reflections are disabled (66.67 percent of the votes, $t = 3.99$, $p < 0.005$);

We also found that reflections were deemed significant in increasing perceived realism of the generated try-on sequences. Videos with reflec-

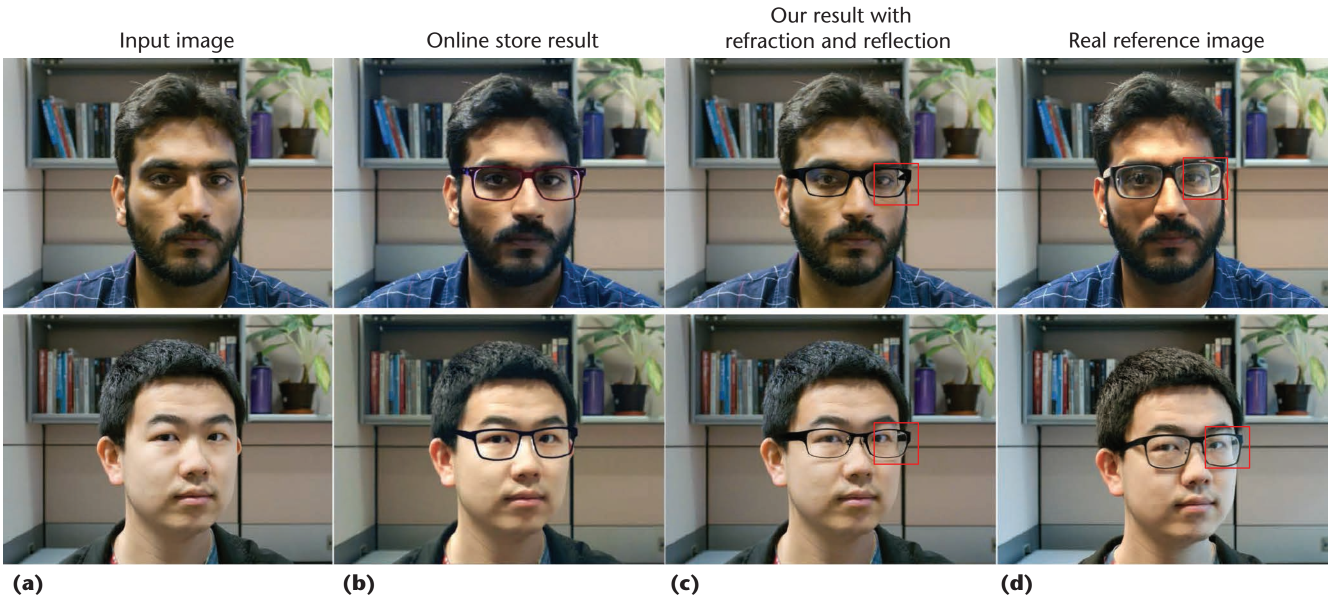

Figure 8. Comparison of eyeglasses virtual try-on results. Given (a) the input image, (b) the virtual try-on solution from an online eyeglasses store only inserts the eyeglasses frame, without considering the lenses. In contrast, (c) our virtual try-on result accounts for the effects of refraction, reflection, and shadows. Our result appears more similar to (d) the real reference image, captured with a similar pose, similar eyeglasses frame, and the same prescription. In both (c) synthetic and (d) real images, the eyes seem smaller due to prescription eyeglasses, and a discontinuity appears along the silhouette of the wearer's face (highlighted with red rectangles).

tions were consistently favored over videos without reflections, both in the videos with refractions (86.67 percent of the votes, $t = 16.68$, $p < 0.005$) and without refractions (89.17 percent of the votes, $t = 19.48$, $p < 0.005$).

Lastly, all our results were systematically favored (greater than 97 percent of the votes) over those of the commercial online virtual try-on. This does not come as a surprise because all our renderings account for the surrounding lighting and exhibit convincing shadows on the face.

Furthermore, comments provided by some participants indicate that once they understood what the differences between stimuli were (refraction and

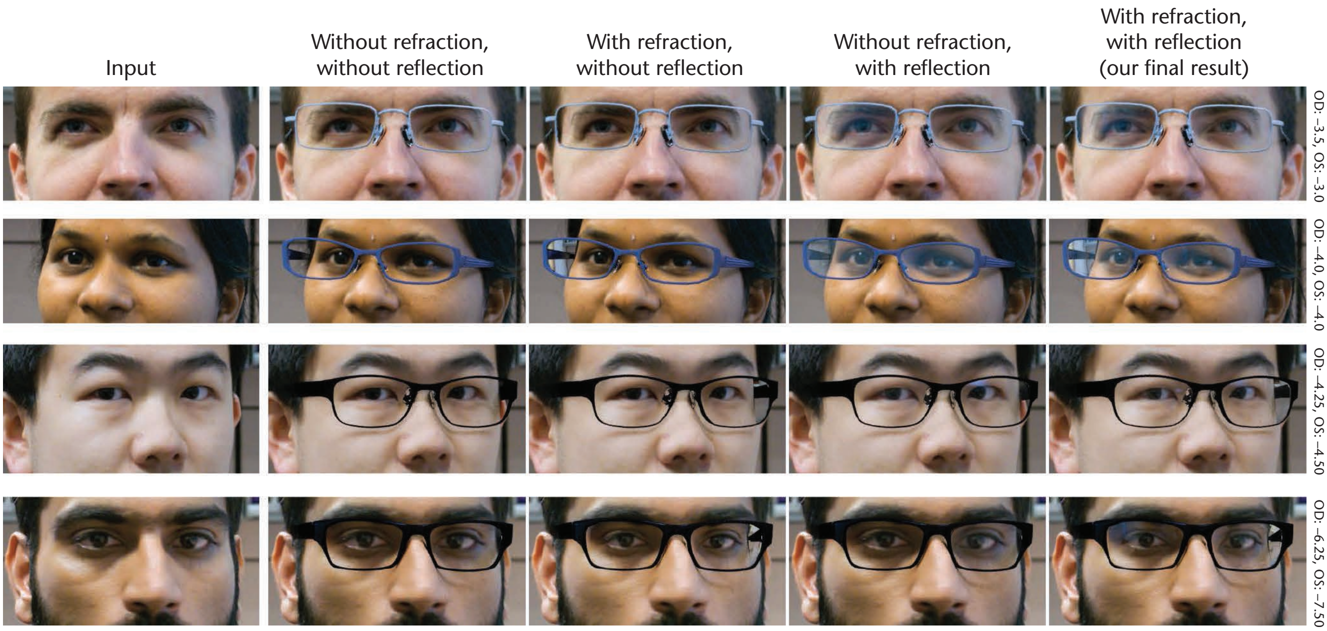

Figure 9. Comparison of four variations of the virtual try-on results generated using our approach, with/without refraction and with/without reflections. The right-hand column shows our final result, with refraction and reflection. The optical power of the prescription lenses are list on the right side of the images.

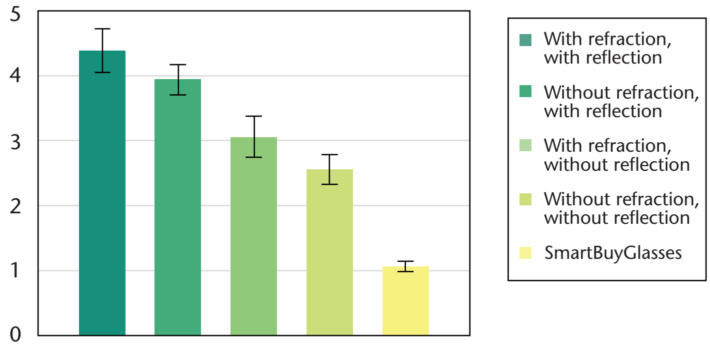

Figure 10. User study results. Bar height indicates the average score for each type of video across all users. Error bars are between-subjects standard error of the mean. Five types of videos were displayed, including four variations of virtual try-on videos generated by our method and one screen capture from a commercial online eyeglasses store (SmartBuyGlasses). The highest-ranking video was assigned a score of 5, and the lowest-ranking video was assigned a score of 1.

reflection), they were able to rank them quickly and consistently.

The user study results support the observation that our eyeglasses virtual try-on system, which takes into account refraction, reflection, and shadows, creates a more realistic experience compared with existing solutions.

Performance

We render the videos on a per-frame basis. The input videos are resized to 720p, with 30 frames per second. For this work, we used a 3.6-GHz Intel Core i7 CPU, and our unoptimized program ran on a Linux virtual machine. On average, it took 5 minutes to render a frame with prescription eyeglasses inserted. Two key parameters that influenced the running time were samples per pixel for ray tracing and samples in hemisphere for Monte Carol integration. For results in this article, we used four (2 × 2) samples per pixel and 160 stratified samples in each hemisphere. Our current rendering of prescription eyeglasses is offline, but the process could be greatly accelerated by multiprocessing or even rendering them in real time using a GPU.

Although the pipeline for our virtual try-on system involves some manual interaction, for example, in the eyeglasses positioning and lens mounting step and the incident lighting capturing step, which takes 12 minutes in total, it could be automated. Eyeglasses can be positioned using affine transformations computed based on preselected feature points on the face and a 3D eyeglasses model.[3] Lenses could also be mounted automatically by aligning with the optical axis of the eyeglasses frame and fitting the lens center into the lens plane of the frame. The incident lighting could be estimated in real-time using spherical harmonics.[23]

We currently insert the eyeglasses based on face tracking. Although individual frames are well rendered, errors in pose estimation could result in wobbling eyeglasses, especially when people turn their heads quickly. This could be alleviated by smoothed head poses. Mapping the rendering to the GPU would let the system function in real time. Future work might employ those alternative techniques to develop a robust real-time system.

Acknowledgments

This research is supported by the BeingThere Center, a collaboration between the Nanyang Technological University Singapore, Eidgenössische Technische Hochschule Zürich, and University of North Carolina at Chapel Hill. The BeingThere Center is supported by the Singapore National Research Foundation under its International Research Center at Singapore Funding Initiative and administered by the Interactive Digital Media Program Office.

References

1. T. Weise et al., "Realtime Performance-Based Facial Animation," *ACM Trans. Graphics*, vol. 30, no. 4, 2011, article no. 77.
2. J.Y. Bouguet, "Camera Calibration Toolbox for

Table 2. Voting results for virtual try-on videos.*

Preferences	With refraction, with reflection	Without refraction, with reflection	With refraction, without reflection	Without refraction, without reflection	Commercial system
With refraction, with reflection	–	62.50%	86.67%	87.50%	99.17%
Without refraction, with reflection	–	–	74.17%	89.17%	100.00%
With refraction, without reflection	–	–	–	66.67%	97.50%
Without refraction, without reflection	–	–	–	–	98.33%
SmartBuyGlasses	–	–	–	–	–

**In the user study, we presented four variations of virtual try-on videos generated by our method and one screen capture from a commercial online eyeglasses store. The percentage represents preference of the row video over the column video—for example, 62.50 percent of the participants favored the "with refraction, with reflection" video over the "without refraction, with reflection" video.*

Matlab," Oct. 2015; www.vision.caltech.edu/bouguetj/calib_doc/.

3. A. Niswar, I.R. Khan, and F. Farbiz, "Virtual Try-On of Eyeglasses Using 3D Model of the Head," *Proc. Int'l Conf. Virtual Reality Continuum and Its Applications in Industry*, 2011.
4. T. Popa et al., *Virtual Mirror Systems and Methods*, patent WO 2015172229 A1, 2015; www.google.com/patents/WO2015172229A1?cl=en.
5. D. Meister and J. E. Sheedy, *Introduction to Ophthalmic Optics*, Carl Zeiss Vision, 2000.
6. P. Debevec, "Rendering Synthetic Objects into Real Scenes: Bridging Traditional and Image-Based Graphics with Global Illumination and High Dynamic Range Photography," *Proc. Siggraph*, 2008, article no. 32.
7. C. Wu et al., "Shading-Based Dynamic Shape Refinement from Multi-view Video under General Illumination," *Proc. Int'l Conf. Computer Vision* (ICCV), 2011, pp. 1108–1115.
8. L. Valgaerts et al., "Lightweight Binocular Facial Performance Capture under Uncontrolled Lighting," *ACM Trans. Graphics*, vol. 31, no. 6, 2012, article no. 187.
9. J. Thies et al., "Real-Time Expression Transfer for Facial Reenactment," *ACM Trans. Graphics*, vol. 34, no. 6, 2015, article no. 183.
10. C. Wu et al., "Real-Time Shading-Based Refinement for Consumer Depth Cameras," *ACM Trans. Graphics*, vol. 33, no. 6, 2014, article no. 200.
11. F. Hill and S. Kelley, *Computer Graphics Using OpenGL*, 3rd ed., Pearson, 2007.
12. T. Whitted, "An Improved Illumination Model for Shaded Display," *Comm. ACM*, vol. 23, no. 6, 1980, pp. 343–349.
13. A.S. Glassner, ed., *An Introduction to Ray Tracing*, Academic Press, 1989.
14. E. Lafortune, "Mathematical Models and Monte Carlo Algorithms for Physically Based Rendering," PhD dissertation, Dept. of Computer Science, Katholieke Universiteit Leuven, 1996.

Qian Zhang *is a research assistant at Nanyang Technological University, Singapore. Her research interests include image processing, computational photography, and image-based rendering. Zhang has a BS from Huazhong University of Science and Technology. Contact her at zhangqian@ntu.edu.sg.*

Yu Guo *is a PhD student in the Department of Computer Science at the University of California, Irvine. During this research, he was a research associate at Nanyang Technological University. His research interests include computer graphics and computer vision, especially geometry and rendering. Guo has an MS in computer science from Shenzhen Institute of Advanced Technology. Contact him at guo.yu@uci.edu.*

Pierre-Yves Laffont *is the CEO and cofounder of Lemnis Technologies. During this research, he was a postdoctoral researcher at ETH Zurich and a visiting researcher at Nanyang Technological University. His research interests include intrinsic image decomposition, example-based appearance transfer, and image-based rendering/relighting. Laffont has a PhD in computer science from Inria Sophia-Antipolis. Contact him at contact@py-laffont.info.*

Tobias Martin *is a research engineer at VirtaMed. During this research, he was a postdoctoral researcher at ETH Zurich. His research interests include geometric modeling, physically based simulation, and rendering. Martin has a PhD in computer science from the University of Utah. Contact him at tobias.martin@virtamed.com.*

Markus Gross *is a professor of computer science at ETH Zurich, head of the Computer Graphics Laboratory, and director of Disney Research, Zurich. Gross has a PhD in computer graphics and image analysis from Saarland University. He is an ACM fellow, received the Technical Achievement Award from the Academy of Motion Picture Arts and Sciences, and has cofounded various startups. Contact him at grossm@inf.ethz.ch.*