EyeContact: Scleral Coil Eye Tracking for Virtual Reality

Eric Whitmire[1], Laura Trutoiu[2], Robert Cavin[2], David Perek[2], Brian Scally[2], James O. Phillips[3], Shwetak Patel[1]

[1]Computer Science and Engineering
University of Washington
{emwhit, shwetak}@uw.edu

[2]Oculus & Facebook

[3]Department of Otolaryngology
University of Washington

ABSTRACT

Eye tracking is a technology of growing importance for mobile and wearable systems, particularly for newly emerging virtual and augmented reality applications (VR and AR). Current eye tracking solutions for wearable AR and VR headsets rely on optical tracking and achieve a typical accuracy of 0.5° to 1°. We investigate a high temporal and spatial resolution eye tracking system based on magnetic tracking using scleral search coils. This technique has historically relied on large generator coils several meters in diameter or requires a restraint for the user's head. We propose a wearable scleral search coil tracking system that allows the user to walk around, and eliminates the need for a head restraint or room-sized coils. Our technique involves a unique placement of generator coils as well as a new calibration approach that accounts for the less uniform magnetic field created by the smaller coils. Using this technique, we can estimate the orientation of the eye with a mean calibrated accuracy of 0.094°.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords

eye tracking; magnetic tracking; scleral search coils; head-mounted displays; virtual reality

INTRODUCTION

Accurate and high-speed eye tracking is important for enabling key scenarios in virtual reality (VR) and augmented reality (AR). Eye tracking could enable a new class of gaze mediated input [7] and techniques such as foveated rendering [5], which can reduce the computational demands of AR/VR by focusing render quality at the users gaze location. Virtual avatars could be made more realistic by including eye tracking information [4], which is impossible to measure with traditional motion capture systems while wearing an head-mounted display (HMD). High accuracy eye tracking could also enable studies of how the human vestibulo-ocular system responds to virtual reality.

ISWC '16, September 12-16 2016, Heidelberg, Germany

ISBN 978-1-4503-4460-9/16/09...$15.00

DOI: http://dx.doi.org/10.1145/2971763.2971771

Figure 1. The EyeContact scleral coil tracker can clip to an HMD and does not require the use of a head mount or room-sized field coils

Existing research on wearable eye tracking systems has focused predominantly on the use and improvement of optical tracking techniques. However, the gold standard for high resolution eye tracking is still magnetic tracking with scleral search coils (SSC) [6]. Scleral coil tracking can record small amplitude motions with high temporal ($> 1\,\text{kHz}$) and spatial resolution (calibrated error $< 0.1°$). In this technique, the head is positioned between large Helmholtz coils, which generate a uniform magnetic field. A wire loop embedded in a silicon annulus is placed on the sclera of the eye. The magnetic field induces a voltage in the scleral coil according to its orientation [13]. By examining the magnitude of the voltages induced in the thin wires leading from the coil, the system estimates the eye's orientation.

One of the major limitations of SSC tracking is the need for large generator coils several meters in diameter or a head restraint such as a bite bar or chin rest [14]. To overcome these limitations, we propose a wearable scleral search coil tracking system compatible with an HMD, such as those used in virtual reality systems. By mounting smaller generator coils directly on the HMD, as shown in Figure 1, we constrain the position of the coils relative to the head, allowing the subject to move freely and eliminating the need for a stationary head or room-sized coils.

The insertion of a scleral search coil is an inherently invasive procedure, typically done with a topical anesthetic, and requires supervision by a trained technician. Hence, we do not recommend the use of a scleral coil tracker for consumer use in VR/AR systems. We instead envision researchers using it when they need eye tracking with high accuracy and high temporal and spatial resolution for wearable HMD scenarios. For example, a psychophysics researcher may use our system to obtain high-resolution data to study microsaccades, small (0.2°) eye movements that occur during fixation [9], while presenting visual targets on the HMD to better understand eye behavior variations. Existing video-based eye trackers for HMDs often have insufficient accuracy and temporal resolution for these kinds of studies. We also hope to encourage researchers to consider our SSC system to obtain ground truth data for evaluating more traditional optical eye tracking systems. An additional possible use for our system is in the medical field, where scleral coil tracking is the gold standard for the diagnosis of subtle vestibular, ophthalmological, and neurological disorders [6].

The specific contributions of this paper are:

1. a mobile, head-mounted system for high-speed and high-accuracy eye tracking that does not require instrumentation of the external environment,
2. a unique coil placement that enables reconstruction of gaze and the position of the scleral coil in space,
3. a calibration technique that accounts for the diverging fields created by the smaller generator coils,
4. a highly accurate mechanical test rig with five degrees of freedom that allows us to thoroughly characterize the magnetic field and enables an evaluation of the accuracy and precision of a SSC tracking system independent of the fixation accuracy of a user.

BACKGROUND AND RELATED WORK

In this section, we discuss the operating principles and research related to magnetic eye trackers and SSC systems. We also briefly discuss alternative methods to magnetic eye tracking (video-based, electrooculography).

Eye Tracking Systems

Video-based tracking is the most widely used method for eye tracking. A high end commercial video eye tracking system, such as the SR Research EyeLink 1000 Plus, is capable of sampling at 1000 Hz with 0.33° average accuracy. Commercial wearable eye tracking systems from Tobii and SMI, take the form of lightweight glasses with IR illumination. These systems often have an accuracy on the order of 0.5° with an output data rate of 60 Hz to 100 Hz but are not designed to be compatible with an HMD. Dual Purkinje tracking is an alternative high precision optical tracking technique, but it requires a complex set of mirrors and servo motors that make it difficult to implement in head-mounted displays. For a more thorough review of wearable video-based eye trackers, see the review by Bulling and colleagues [2].

Video-based eye tracking solutions have recently been adapted for use in an HMD. For example, the SMI tracking system is available as an add-on package for the Oculus Rift DK2 and Samsung Gear VR. It achieves 0.5° to 1° accuracy with a data rate of 60 Hz. Additional solutions from Arrington Research and Tobii offer similar performance. Although these solutions will likely improve with time, current solutions suffer from a low data rate and accuracy and are designed specifically for the optics of a particular HMD.

Another older method for eye-tracking is electrooculography (EOG). The eye is the source of an electric dipole between the cornea and the retina; this corneoretinal potential can be up to 1 mV when recorded with EOG and varies as a function of the orientation of the eye. Electrooculography measures eye movement by measuring this corneoretinal potential through electrodes placed around the eye. EOG signals are noisy and are often only used for horizontal eye motion or to detect blinks and eye gestures [2].

Scleral Coil Tracking

Scleral coil tracking is the gold standard for eye tracking, particularly within the medical research community. The technique was developed by Robinson in 1963 [13], refined by Collewijn and colleagues in 1975 [3], and offers a data rate often approaching 10 kHz with a calibrated accuracy better than 0.1°. The subject wears a wire coil embedded within a silicone annulus that sits on the sclera of the eye (Figure 2). A thin wire connects to an external measurement unit. External generator coils create a magnetic field that induces a voltage in the coils, which is amplified and measured.

Theory of Operation for Scleral Coil Trackers

Scleral coil tracking relies on alternating magnetic fields present at the user's eyes. Using an alternating current to drive an electromagnet formed by a wire coil is an effective method of generating a magnetic field that oscillates at a particular frequency. According to the Biot–Savart law, an electric current flowing along a wire coil will generate a magnetic field that resembles Figure 3.

The oscillating magnetic flux from these generator coils intersects the scleral contact and induces a current of the same frequency within the scleral coil according to Faraday's law of induction. The current induced in the scleral coil is proportional to the rate of change of the magnetic flux through the surface bounded by the coil.

The flux through the coil depends on the orientation of the coil within the magnetic field. If the coil is aligned with the field (that is, the normal vector to the coil is aligned with the field), then the flux and the magnitude of the induced voltage will be maximal. As the coil is rotated away from the field, the induced voltage decreases to zero. If the coil is flipped around, the voltage will have a 180° phase shift, which can be represented as a negative amplitude.

Using this phenomenon, SSC trackers estimate eye orientation by measuring the magnitude of the voltage signal in the scleral

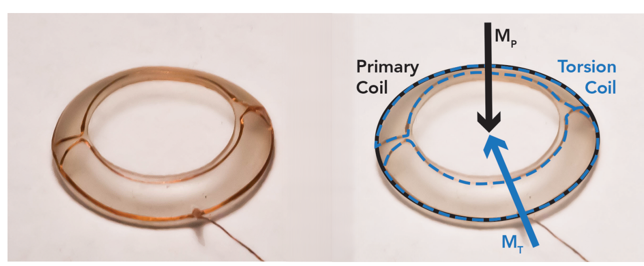

Figure 2. 3D torsional scleral coil from Chronos Vision. The primary coil captures flux flowing through the contact while the torsional coil captures flux traveling across the contact.

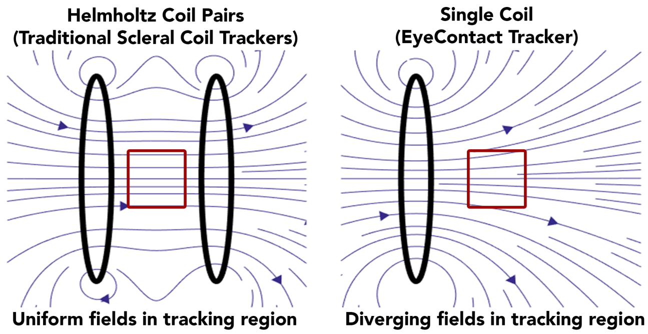

Figure 3. (left) Traditional scleral coil trackers generate fields using Helmholtz coil pairs, which produce a small volume where the magnetic field is insensitive to small head movements. (Right) Fields generated by the unpaired coils in our system create a diverging magnetic field.

coil. However, since the information from a single magnetic field is not enough to completely specify gaze, most scleral coil trackers use three orthogonal generator coils operating at different frequencies or in quadrature. By examining the scleral coil signal at different frequencies, one can estimate the magnetic flux due to each of the generator coils. Traditionally, the magnetic fields are generated by large coil pairs known as Helmholtz coils [13]. As shown in Figure 3 (left), these coil pairs simplify gaze estimation by creating a region with uniform magnetic fields in the x, y, and z-directions. Within that region, the magnetic field direction and magnitude are relatively insensitive to the scleral coil position, allowing users a small amount of movement. With smaller Helmholtz coil systems (less than 1 m in diameter), the uniform volume is only a few centimeters wide [14] and restraints are used to keep the head within this volume. These design limitations impose significant space and cost constraints, preventing the adoption of this technique outside of dedicated facilities in medical centers. Today, researchers use scleral coil eye tracking for the diagnosis of vestibular disorders and biomedical research.

Some SSC systems also seek to measure torsion, which refers to rotation of the eye about its visual axis. With traditional coil systems, a single coil within the scleral contact is not sufficient to determine the torsion of the eye. Consequently, some scleral contacts contain a second, torsional coil. As shown in Figure 2, this coil forms two loops around the eye on two different planes. While the primary coil captures flux through the contact, any current induced in the torsional coil is due to magnetic flux traveling across the coil. It is conceptually convenient to consider the torsional coil as a second coil perpendicular to the primary coil.

Improvements to Scleral Coil Tracking

In an attempt to increase the usability of SSC methods, several projects have sought to eliminate the wire connecting the scleral coil to the measurement device. Reulen and Bakker describe a double magnetic induction (DMI) technique [11] that uses a short-circuited scleral coil. As before, current is induced in the scleral coil loop when the user is placed within an alternating magnetic field generated by a set of Helmholtz coils. Current within this loop induces a secondary magnetic field, which can be detected by a set of coils directly in front of the eye. Similar work by Bremen and colleagues [1] showed a wearable version of the DMI technique. Fundamentally, DMI suffers from a weak signal to noise ratio because of the need to measure the secondary magnetic field. Moreover, they still require large external coils to generate magnetic fields. In our EyeContact system, we use the standard wired SSC and focus on innovation on the generator coil design. As a result, our system is comparable in accuracy to standard SSC trackers, does not require instrumentation of the external environment, and can be attached to an HMD.

Roberts and colleagues showed a wireless scleral coil system that operates on a similar principle [12]. By shorting the scleral coil with a series capacitor embedded in the silicon annulus, they created a resonant circuit. The resonant circuit produced an oscillating current in the coil, sustained by power fed by a nearby transmitter coil driver. Several receiver coils measured the decaying signals. The relative strength of the signal in the receiver coils determines the orientation of eye gaze. Although it may be possible to adapt for use with an HMD, the performance specifications of such a technique have not been verified.

Thomassen and colleagues noted that the small uniform magnetic field region of many trackers made it difficult to carry out head-unrestrained experiments [14]. The authors adapted the tracking algorithm of an off-the-shelf scleral coil tracking system to correct its performance when the head moves out of the uniform region. However, this technique still requires significant instrumentation of the environment. Plotkin and colleagues also attempted to address the nonuniformity problem with a planar transmitter placed in a fixed location behind the user [10]. Although their design is more portable, it is too large to be mounted on the user, and does not allow the user to move more than 10 cm during an experiment. In contrast, our system allows the user to move around freely, limited only by the tether for the HMD, scleral coil, and field coils.

SYSTEM DESIGN

We present EyeContact, a scleral coil tracking system designed for use with virtual reality headsets. We use five small coils as shown in Figure 4 (left), and rigidly mount the coils on the user's head, constraining the magnetic flux orientation to the user's head orientation. Each coil oscillates at a unique frequency and creates a unique magnetic field. A key challenge in our system is that, in contrast to Helmholtz coils, these coils generate diverging fields and the field orientation strongly varies with position, as shown in Figure 3 (right). Further complicating our system is the fact that a scleral coil translates in space along the surface of the eye as the eye rotates. As it moves relative to the field coils, the magnitude and direction of the magnetic fields at the scleral coil change. However, since the decomposed five magnetic fields are unique at each position in space, gaze estimation is still possible. Because the fields vary with eye position, we are also able to recover the positional offset between the tracker and the scleral coil. This allows us to account for any shifts or slippage of the HMD on the user's face.

The use of five generator coils in this arrangement ensures that three coils are close to each eye and can provide a reliable signal (the central generator coil is shared by both eyes).

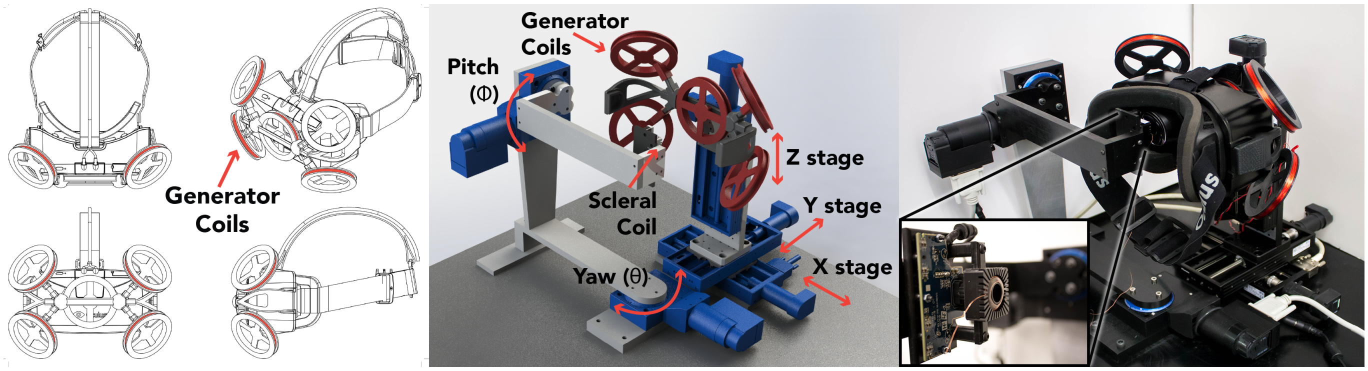

Figure 4. (left) The scleral coil tracker consists of five field coils arranged around the HMD. It is designed to snap on to the Oculus Rift DK2, but it can be easily adapted for other HMDs. (center, right) The 5-DOF mechanical evaluation rig allows us to control the scleral coil orientation (pitch, yaw) and adjust the position of the HMD/tracker with respect to the scleral coil. (right, inset) The scleral coil rests in the arm of the test rig.

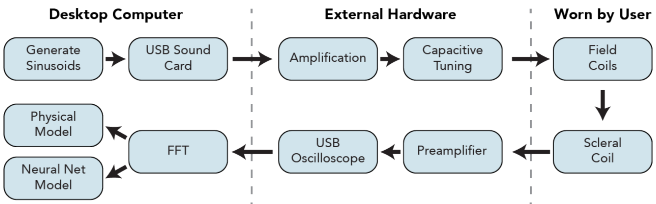

Figure 5. Scleral coil tracking signal pipeline. Sinusoids are synthesized from a desktop computer and amplified before passing through the generator coils. This induces a voltage in the scleral coil, which is amplified and processed on a desktop computer.

Although these three generator coils alone are sufficient for reconstructing the orientation of the eye, the two coils furthest from each eye provide a weak signal that further improves the accuracy of the gaze estimation. The size of the coils was optimized to balance magnetic field strength and the weight of the tracking device. By placing the coils close to the eyes, we can generate magnetic fields with comparable strength to traditional field coils while using much less current (less than 1 A per coil).

For the scleral coil, we use a 3D scleral contact (which contains a torsional coil) from Chronos Vision. In contrast to previous efforts, we use the torsional coil for more than just torsion estimation. Rather, we consider the entire scleral coil as a biaxial magnetic field sensor that measures two of the three components of the magnetic field.

Each generator coil consists of 50 turns of 26 AWG magnet wire wound around a circular 3D printed ABS frame, 8 cm in diameter. The coils fit together and form an attachment for an HMD. Our implementation is designed to fit on an Oculus Rift DK2, but could easily be modified to fit other VR devices. A system diagram is shown in Figure 5. The coils are driven by sinusoidal currents synthesized from a desktop computer. By keeping the frequency of the coil stimulus in the usable audio range, commodity audio hardware can be used for synthesis and amplification. We use a desktop computer running Max 7 software to synthesize sinusoids with a 192 kHz sampling rate and frequencies of approximately 15 kHz, 16 kHz, 17.1 kHz, 18.3 kHz, and 19.6 kHz. The frequencies were chosen to align with Fast Fourier Transform (FFT) bin locations and were spaced unequally to avoid intermodulation components.

The synthesized audio signals are converted to analog voltage signals using an 8-channel USB audio interface. The five voltage signals from the audio interface are amplified with independent class D amplifiers to increase the current through the field coils.

Because the field coils present an inductive load to the system, a capacitive tuning adapter is used to remove the imaginary component of the impedance. The tuning adapter consists of five channels, each with a step-up transformer and a bank of parallel capacitors. By tuning the capacitance of each channel while monitoring the voltage and current through the coil, it is possible to maximize the power transfer to the coils.

Prior to measuring the signals from the scleral coil, we amplify them using an instrumentation amplifier (INA128).We then sample the primary and torsional signals using a 4-channel differential input USB oscilloscope from Pico Technology.

Software running in MATLAB on a desktop computer interfaces with the USB oscilloscope and records data at 1 MSa/s. The software buffers the signal for each eye from the primary and torsional coil into 16 ms buffers with 25% overlap. A 4-term Blackman-Harris window is applied to each buffer, since their lengths are not an integer number of periods of the field frequencies. This window function was selected because of its side-lobe attenuation. We then compute the FFT for each window and select the five primary and five torsional components corresponding to the frequency bins of interest. By using only these narrow frequency bins, we also avoid any magnetic interference from the environment or the electronics of the HMD; the scleral coil measurements are not changed when the HMD is powered on or an application is in use. For each frequency component, we save the magnitude and phase of the complex FFT. Optionally, sequential measurements can be averaged to further improve the signal-to-noise ratio.

In selecting frequencies that align with the FFT samples, we also force each window to contain an integer number of periods of that particular frequency, maintaining the phase of the FFT between windows. However, due to small timing differences between the components of the system, in practice, the phase drifts significantly over time. To account for this drift, we use a 10 second calibration period during which we measure and track the phase drift rate. Periodically, we reset any additional drift that has accumulated. With this technique, the adjusted phase shifts only occur in 180° increments when magnetic flux through the coil is reversed due to a change in gaze orientation. We can then reconstruct the signed magnitude by comparing the adjusted phase with the phase at a known gaze orientation.

MECHANICAL TEST RIG

To evaluate the performance of our scleral coil tracker, we constructed a mechanical 5-DOF test rig, capable of adjusting the orientation of the coil and the position of the HMD with respect to the coil as shown in Figure 4 (center, right). The scleral coil is placed inside a plastic holder mounted on a U-shaped arm. Motorized rotation stages from Physik Instrumente allow the arm to swing around (θ) and to tip up and down (ϕ). This allows us to adjust the yaw (θ) and pitch (ϕ) of the coil direction vector with a resolution of 0.0018°. The HMD and generator coils are mounted on a three-axis linear stage system (x, y, z). This allows us to move the HMD in space with respect to the coils with a resolution of 0.5 µm, simulating different placements of the HMD on the forehead or slippage of the HMD.

Because we adopt a spherical model of human eye, a scleral contact placed on the surface of the eye moves in space as the eye rotates, according to the radius of the eye (R). For a system with a non-uniform magnetic field, this is an important consideration, as the magnetic field environment will vary with respect to gaze, even if the HMD is fixed to the user's head. To account for this, the base of the scleral contact is mounted $R_e = 8.5$ mm in front of the center of rotation, simulating the behavior of a human eye. To compute the position offset of the scleral coil, we consider both the offset due to the HMD position (x, y, z) and due to the orientation of the eye (θ, ϕ) according to Equation 1.

$$P = \langle x - R_e \sin(\theta)\cos(\phi), y - R_e \cos(\theta)\cos(\phi), z - R_e \sin(\phi)\rangle \quad (1)$$

The five motorized stages can be controlled from a desktop computer. We have integrated control of the stages into our MATLAB data collection software to enable an automated sweep of the entire visual field. When measuring the fields, we sweep the coil in the desired range in yaw (θ) and pitch (ϕ) and sweep the HMD position in space (x, y, z). The MATLAB recording software automatically records the scleral coil values one second after stage movement has ceased, to avoid any effects from vibration of the test rig.

Data is collected from $-30°$ to 30° in 3° increments along the horizontal (θ) and vertical (ϕ) axes. This is comparable to the operating range of other eye tracking devices. We simulate slippage by sweeping the HMD from -7.5 mm to 7.5 mm in 3 mm increments along the x-axis (side to side), -5 mm to 0 mm in 5 mm increments along the y-axis (front to back), and -5 mm to 5 mm in 2 mm increments along the z-axis (up and down), as shown in Figure 4 (center). This represents a full sweep of 441 points at $6 \times 2 \times 6 = 72$ locations in space for a total of 31 752 data points. Collecting this data takes approximately 8 h.

GAZE ESTIMATION MODELS

We compare two models for gaze estimation. First, we use a physics approach that models the magnetic fields around the tracker. By comparing the scleral coil measurements with the values we would expect to measure given the magnetic fields, we can estimate the scleral coil position and orientation. We compare this approach to a neural network model that directly estimates orientation given the scleral coil measurements.

Gaze Estimation using Physical Model

The physical model estimates the five magnetic field directions at the eye in world space and then uses the scleral coil measurements to reconstruct an estimate of those vectors in coil space. The eye orientation is determined by the rotation that best accounts for this coordinate system transform. As summarized in Figure 6, the model consists of a per-use calibration to estimate the offset of the tracker with respect to the eye and a per-frame gaze orientation estimation procedure.

The five magnetic fields are modeled as wire loops at fixed positions and orientations. Based on the known geometry of the tracker and the current through each generator coil, we can compute the expected magnetic field, $B(P)$, at any potential eye location, according to standard equations for the magnetic field of a coil off the symmetric axis. With this information, given the position and orientation of the scleral coil, we can predict the values we would expect to measure from the scleral coil ($\hat{M}(B, P, \theta, \phi)$).

Calibration Procedure

When the user first puts on the HMD, there is an unknown offset between the tracker and the eye's center of rotation due to differences in user anatomy and HMD placement. Our calibration procedure seeks to estimate this HMD offset. During the procedure, we must also learn the sensitivity of the primary and torsional coils as well as the radius of the eye, as this determines how much the scleral coil moves as the eye rotates. We use a 36-point calibration procedure in which the user looks at targets placed at specific locations on screen (see Figure 6, top-left). We simulate this in the mechanical test rig by moving the stages to the required orientations.

Given these calibration points (with a known gaze orientation) and the magnetic field models, we use MATLAB's Global Optimization Toolbox to estimate the scleral coil sensitivities, the offset of the HMD (x_e, y_e, z_e), and the eye radius (R_e). We evaluated our ability to estimate the HMD offset by using the test rig to shift the HMD and calibrating at 72 different offsets within the slippage volume. At each point, we used 36 points for the calibration procedure. Mean Euclidean error in estimating the HMD offset was 0.72 mm ($\sigma = 0.24$ mm).

This model reasonably explains the measurements observed in the scleral coil (M_{meas}). However, there is a small error (mean error for primary coil was 1.5%) in these estimations, caused by not accounting for irregularities in eye movement, irregularities in the magnetic fields due to physical construction, and distortions in the field caused by the presence of the HMD. An important realization, however, is that these errors are systematic. It is possible to model the residuals between each of the ten scleral coil readings and those predicted by

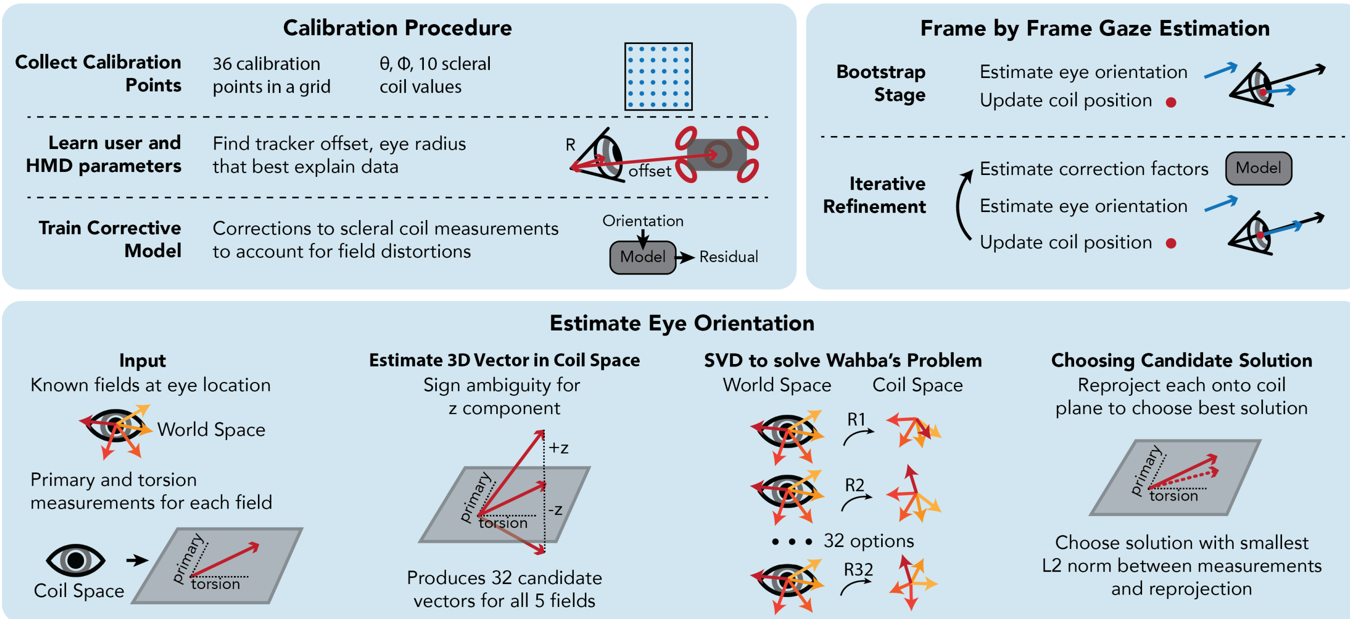

Figure 6. Summary of the calibration and gaze estimation procedures. A 36-point calibration procedure is performed when the user puts on the HMD to estimate the HMD offset and eye radius and to train a corrective model for the scleral measurements. Each frame, eye orientation is iteratively estimated and refined using updated coil positions and correction factors. Eye orientation is estimated using an SVD solution to Wahba's problem.

the physics model ($M_{meas} - \hat{M}$) as a function of gaze location. That is, if we knew the user's gaze, we could adjust the measurements from the scleral coil such that they would best fit the predicted values from the physics model. When modeling the residuals as a function of gaze position, it is possible to learn the corrective model using just the 36 calibration points. We use MATLAB to fit a function ($F(g_x, g_z) = \hat{M} - M_{meas}$) to the 36 calibration points using biharmonic interpolation, for each of the 10 scleral coil measurements. During gaze estimation, we rely on an iterative bootstrapping approach to refine the gaze estimate using the corrective models.

Frame by Frame Gaze Estimation

When estimating the orientation of the eye, we first use the magnetic field model to compute the five expected magnetic field directions at the scleral coil position, which can be calculated using the HMD offset learned during the calibration. Initially, since we do not know the eye orientation, we do not account for the traversal of the coil in space as the eye moves; we assume, for now, a fixed position of the coil. Next we reconstruct the five measured field directions in coil space from the scleral coil measurements. We consider the scleral coil as biaxial sensor that outputs the projection of the magnetic field vector in the plane of the scleral coil, which is determined by the orientation of the primary and torsion coils (see Figure 6, bottom). To reconstruct the third component, we normalize the measurements by the field strengths and by the relative sensitivities of the primary and torsional coils, so that the scleral coil sensor effectively reports two components of a unit vector. We can reconstruct this third component as $M_z = \pm\sqrt{1 - M_p^2 - M_t^2}$. Note that there is a sign ambiguity for this third field component. We consider both possibilities for each of the five fields, producing 32 candidate sets of field measurements.

For each candidate set, we now have five known field vectors (from the physics model) and five estimates of those fields (from the scleral coil, with the 3rd component estimated) in the unknown reference frame of the coil. This is an instance of Wahba's problem, which seeks to find a rotation matrix between two coordinate systems given a set of weighted observations. We use an SVD solution to Wahba's problem to estimate the gaze [8]. The weights for each field are set based on a precomputed average signal-to-noise ratio for each coil.

This procedure outputs 32 candidate solutions that fully specify the eye orientation. To choose among the solutions, we compute the expected values of the scleral coil measurements for each of the candidate solutions and choose the solution that most closely matches the measured values.

This initial gaze estimation does not account for the corrective model or the movement of the eye as it rotates. Given the initial gaze estimate, which tends to be 3° off on average, we can compute an estimate for the eye position and necessary corrective factors. To refine our gaze estimation, we compute the updated magnetic field directions at the new eye location, adjust each of our scleral coil measurements according to the corrective model ($M_{adj} = M_{meas} + F(\hat{g}_x, \hat{g}_z)$) and recompute the estimated gaze orientation. This next gaze estimate is much more accurate and allows an even better estimate of the corrective terms and scleral coil position. We repeat this process until the estimated gaze converges; typically this requires 5 iterations. Applying this procedure to the 405 test orientations across all 72 HMD offsets, we achieve a mean error of 0.18°. The model performance for a small subset of these points is depicted in Figure 7 (top right) and a complete

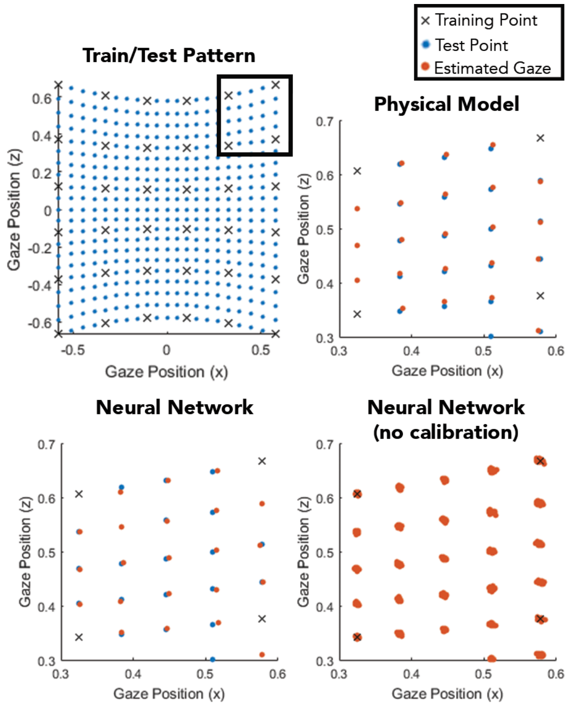

Figure 7. Gaze plot showing training (black X), test (blue circles), and estimated (orange circles) gaze locations as projected onto a plane at $y = 1$. (top left) 36 points were used for training and the remaining 405 used for testing. (bottom and right) The physical and neural network model gaze estimations and ground truth for a small subset of the field of view. (bottom-right) The calibration-free neural network shows gaze position for all 72 HMD offsets.

summary is shown in Figure 8 as a cumulative distribution function (CDF).

Neural Network Model

Rigid Body Model

As an alternative to the physical model, we trained a neural network to learn the effects of field divergence, distortion caused by the HMD, and movement of the eye. To train the model, the user performs a short 36-point calibration procedure. Each data point consists of the 10 scleral coil measurements. Forty neurons were used for the hidden layer and the output layer consists of two neurons, which were trained to output the orientation of the scleral coil (θ and ϕ).

As with the physical model, we used the test rig to evaluate the accuracy of this gaze estimation model. At each of the 72 HMD offsets within the slippage volume, we train a separate network using 36 training point and holding out the remaining 405 points for testing. This approach achieves 0.094° mean error, averaged over all HMD offsets and test points. Gaze estimations for one HMD offset are shown in Figure 7 (bottom left). The full results are summarized as a CDF in Figure 8.

Calibration Free Model

We also sought to explore whether the neural network can be trained in a position-independent manner in order to reduce sensitivity to slippage. In this approach, a two-stage model is used to both estimate the search coil position and orientation. A significant amount of training data is needed to train these models, making a calibration procedure unfeasible. Moreover, it is impossible to position the search coil at a specific location in space while it is being worn by a user. Hence, we introduce a calibration-free tracking technique using a pre-trained model.

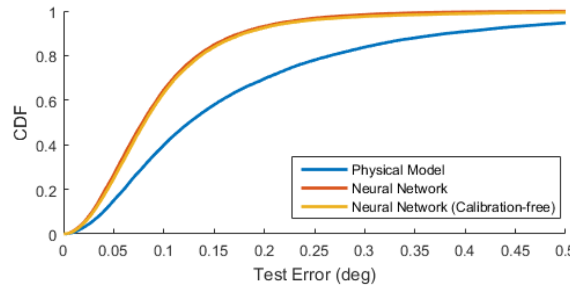

Figure 8. CDF of gaze orientation accuracy across all positions with the physical model, neural network model, and calibration-free neural network model. The physical model achieves $0.18°$ mean error. The neural network models both outperform the physical model with mean errors of $0.094°$ and $0.099°$, respectively.

First, a neural network with 10 input nodes (corresponding to the 5 primary and 5 torsional observations), 40 hidden nodes, and 3 output nodes (x, y, z) is trained using calibration points collected at 32 of the 72 HMD offset positions. Mean Euclidean error on the remaining 27 880 points was 0.084 mm.

We now augment our original gaze model with the position estimation. Input consists of 13 nodes, 10 from the scleral coil and the estimated x, y, z location in space. Again, we train on calibration points collected the training positions and test on the remaining points. Mean error on the test set was 0.099°. These results are summarized in the target plots in Figure 7 (bottom right) and in the CDF in Figure 8.

DISCUSSION

We present two gaze estimation methods based on a physics model and a neural network. The physics model achieves a mean gaze orientation error of 0.18°. By modeling the magnetic fields, this model allows us to understand the magnetic environment of the scleral coil and lends itself well to extensions. For example, future work could look at online calibration procedures, torsion estimation, or using the data from both eyes in calibration. These applications are more difficult with the more opaque neural network model.

With a mean error of 0.094°, the neural network model is better able to approximate the magnetic environment around the tracker and the distortions caused by the HMD. The calibration-free neural network model can significantly improve usability by estimating both the HMD slippage and eye orientation in each frame. However, since this model was trained on data collected from the test rig, it is unclear how well it generalizes to a human user in a different magnetic environment. Experiments are underway to validate the performance of this technique with human participants.

A promising avenue for future exploration is the combination of these models. By providing a neural network with the magnetic field estimations from the physical model, we can simplify the problem and enable shorter training times.

In the calibration-based models, we expect the user to perform the 36-point calibration procedure when they first put on the

HMD and whenever the HMD shifts on their face. We have implemented a shift detection algorithm that can automatically prompt the user to recalibrate, if needed.

Our system samples the scleral coil signal at 1 MSa/s. After buffering and windowing, the gaze estimation is output at a rate of 244 Hz. Significant increases in output data rate can be obtained by reducing the window size or increasing overlap width. We intend to implement a native code version of the algorithm or use special purpose programmable logic hardware. Similarly, we can increase temporal resolution by using higher frequency sinusoids and taking advantage of additional bandwidth. We also explored temporally averaging samples together, as is common in many SSC tracking systems. Though we can improve the accuracy on the neural network models to 0.03°, we chose maintain a higher data rate instead.

In our evaluation, we focused on the mechanical test rig as a proxy for a human user. This allows us to collect a comprehensive dataset of positions and orientations and gives us reliable ground truth references that can isolate the accuracy of the tracker, independent of the fixation accuracy of a human user. In designing the test rig, we accounted for many nuances that make eye tracking in an HMD difficult. For example, we accounted for movement of the scleral coil in space by offsetting the scleral coil location from the center of rotation of the test rig. We also accounted for HMD slippage by mounting the HMD and tracker on a 3-axis translational stage system. Since for many VR applications, measuring the gaze position is sufficient, we did not prioritize evaluating eye torsion, which would provide a full-attitude gaze estimation. Though our physical model provides an estimate of torsion, we do not evaluate its accuracy because the test rig only allows us to adjust the torsion of the scleral coil in 10° increments. A future version of this test rig could include precise control of torsion in the $-4°$ to $4°$ range, which would enable a more rigorous evaluation of the torsion estimation from our model. However, the calibration-based models do account for systematic torsion that occurs regularly as a function of gaze.

CONCLUSION

EyeContact is a scleral coil tracking system designed for a virtual or augmented reality HMD that enables high-speed, high-accuracy mobile eye tracking without instrumentation of the environment. Our mechanical test rig enables calibration and evaluation with a reliable ground truth. We describe a calibration procedure and two different approaches to estimate gaze: a physical model with 0.18° mean error and a neural network model with 0.094° mean error. We hope this system will be a useful tool for researchers in need of high-quality eye tracking within an HMD.

REFERENCES

1. Peter Bremen, Robert F Van der Willigen, and A John Van Opstal. 2007. Applying double magnetic induction to measure two-dimensional head-unrestrained gaze shifts in human subjects. *Journal of neurophysiology* 98, 6 (2007), 3759–3769.
2. Andreas Bulling, Daniel Roggen, and Gerhard Tröster. 2009. Wearable EOG goggles: Seamless sensing and context-awareness in everyday environments. *Journal of Ambient Intelligence and Smart Environments* 1, 2 (2009), 157–171.
3. H Collewijn, F Van der Mark, and TC Jansen. 1975. Precise recording of human eye movements. *Vision research* 15, 3 (1975), 447–IN5.
4. Maia Garau, Mel Slater, Vinoba Vinayagamoorthy, Andrea Brogni, Anthony Steed, and Martina Angela Sasse. 2003. The impact of avatar realism and eye gaze control on perceived quality of communication in a shared immersive virtual environment. In *Proceedings of the 2003 Conference on Human Factors in Computing Systems, CHI 2003, Ft. Lauderdale, Florida, USA, April 5-10, 2003*. 529–536.
5. Brian Guenter, Mark Finch, Steven Drucker, Desney Tan, and John Snyder. 2012. Foveated 3D graphics. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 164.
6. Mark MJ Houben, Janine Goumans, and Johannes van der Steen. 2006. Recording three-dimensional eye movements: scleral search coils versus video oculography. *Investigative Ophthalmology & Visual Science* 47, 1 (2006), 179–187.
7. Kyung-Nam Kim and RS Ramakrishna. 1999. Vision-based eye-gaze tracking for human computer interface. In *Systems, Man, and Cybernetics, 1999. IEEE SMC'99 Conference Proceedings. 1999 IEEE International Conference on*, Vol. 2. IEEE, 324–329.
8. F Landis Markley and Daniele Mortari. 1999. How to estimate attitude from vector observations. (1999).
9. Susana Martinez-Conde, Jorge Otero-Millan, and Stephen L Macknik. 2013. The impact of microsaccades on vision: towards a unified theory of saccadic function. *Nature Reviews Neuroscience* 14, 2 (2013), 83–96.
10. Anton Plotkin, Oren Shafrir, Eugene Paperno, and Daniel M Kaplan. 2010. Magnetic eye tracking: A new approach employing a planar transmitter. *Biomedical Engineering, IEEE Transactions on* 57, 5 (2010), 1209–1215.
11. JPH Reulen and L Bakker. 1982. The measurement of eye movement using double magnetic induction. *Biomedical Engineering, IEEE Transactions on* 11 (1982), 740–744.
12. Dale Roberts, Mark Shelhamer, and Aaron Wong. 2008. A new wireless search-coil system. In *Proceedings of the 2008 symposium on Eye tracking research & applications*. ACM, 197–204.
13. David Robinson and others. 1963. A method of measuring eye movemnent using a scieral search coil in a magnetic field. *Bio-medical Electronics, IEEE Transactions on* 10, 4 (1963), 137–145.
14. Jakob S Thomassen, Giacomo Di Benedetto, and Bernhard JM Hess. 2010. Decoding 3D search coil signals in a non-homogeneous magnetic field. *Vision research* 50, 13 (2010), 1203–1213.