

Real-Time Drawing Assistance through Crowdsourcing

Alex Limpaecher[†], Nicolas Feltman[†], Adrien Treuille[†], Michael Cohen^{*}

[†]Carnegie Mellon University ^{*}Microsoft Research

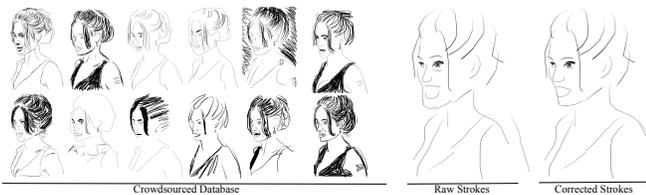


Figure 1: Left: twelve drawings of Angelina Jolie by DrawAFriend players. Right: real-time corrected drawings based on a consensus of the drawings by previous players.

Abstract

We propose a new method for the large-scale collection and analysis of drawings by using a mobile game specifically designed to collect such data. Analyzing this crowdsourced drawing database, we build a spatially varying model of artistic consensus at the stroke level. We then present a surprisingly simple stroke-correction method which uses our artistic consensus model to improve strokes in real-time. Importantly, our auto-corrections run interactively and appear nearly invisible to the user while seamlessly preserving artistic intent. Closing the loop, the game itself serves as a platform for large-scale evaluation of the effectiveness of our stroke correction algorithm.

Drawing dates to well before other forms of recorded history. Today, drawing remains a vital form of artistic expression and an important window into human perception. However, the central challenge to further scientific analysis of drawing is data scarcity. Although search engines index a huge collection of line drawings, these images are stored in raster format with little or no useful metadata. Ideally, a drawing corpus would contain precise stroke-level data for each image, including timing information. We would also like semantic metadata identifying artists and subjects. Even more ambitiously, we would like to glean *perceptual* information, such as which strokes contributed most to image recognition. Finally, for statistical purposes, we would like a *large* dataset, with many drawings by the same artist and many drawings of the same subject by different artists.

Copyright © 2013, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

To address this challenge, we developed *DrawAFriend*, an iPhone game designed to collect drawing data, including all of the information described above. We currently focus on face portraits. Faces are exceedingly difficult to draw by hand, even more so using a touch interface on a small mobile device. To aid users and to collect multiple drawings of the same subject, we allow players to trace over existing photographs. In its first week of release *DrawAFriend* generated over 1,500 images per day.

We believe this large and continuously growing drawing database will enable a rich stream of future research in graphics. As a first application, we demonstrate how the *DrawAFriend* corpus can be mined to provide a self-correcting touch-based drawing interface on mobile devices. We observe that drawing with a touch device often suffers from the “fat finger” problem. We factor this issue into two elements: (1) the “intent” of the artist in drawing a stroke, and (2) an additional random noise component caused by inaccuracy in the touch interface. We therefore hypothesize that if we can determine a consensus of strokes (in an appropriate sense) over a sufficiently large database of drawings, then we can cancel out the noise and recover the artist’s original intent. We analyze the drawing corpus to compute a *correction vector field*, using a modified mean shift algorithm, that for any location, points towards a nearby consensus of strokes. This allows us to develop a real-time *self-correcting* touch interface: as players draw, we essentially clean up their drawings by using data from previous drawings of the same subject. We further introduce a surprisingly simple method to correct strokes based on this consensus while maintaining the *stylistic* choices of the artist. The interface requires no new user interaction paradigms; in other words, it appears “invisible” to the user. The resulting strokes feel more like the intent of the user than the raw original strokes.

To validate the effectiveness of our auto-correction algorithm, we ran a large-scale user study within the game. Each time a user draws a celebrity, we randomly turn the stroke correction on or off. Our results validate the effectiveness of our stroke correction algorithm: with autocorrect on, artists do not need to draw as accurately, and they undo their strokes less. The ability of *DrawAFriend* to serve simultaneously as a large-scale visual data collection platform and as a statistically relevant user study underscores the generality of our crowdsourcing approach.

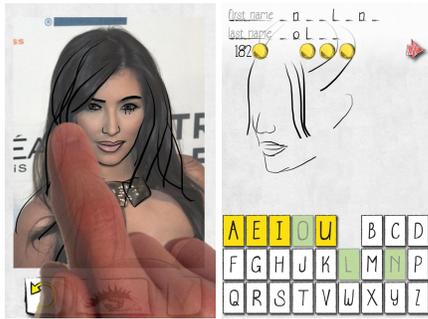


Figure 2: DrawAFriend: tracing a photo (left), guessing identity (right).

DrawAFriend: The Game

To capture and analyze a large-scale drawing dataset, we have developed *DrawAFriend*, a Facebook-integrated turn-based drawing and guessing game for mobile devices. The game is designed to intrinsically motivate players to contribute drawings through a hangman-like guessing mechanism. This approach enables us to gather a large number of drawings with zero marginal cost per drawing, and to modify and instrument the game to capture specific types of data.

The game works as follows. Players have an option to start a game with either a Facebook friend or an anonymous stranger. The player is then given four pictures which she can draw. After choosing a photo to draw, the player can trace the image (see Figure 2 left). Once finished, the player sends her drawing to the friend or anonymous player with whom she is playing. The friend receives a notification that they have a drawing to guess. The user is prompted to guess the identity of the other player’s drawing (Figure 2 right). The drawing is replayed stroke by stroke, and similar to *Hangman*, the player can guess which letters are in the mutual friend or celebrity’s name.

The tracing paradigm results in a set of pre-aligned drawings. Furthermore, by observing the guesses we can indirectly evaluate the quality of the drawings. This dataset includes drawings from artists around the world with different artistic and cultural backgrounds.

Data Driven Drawing

Our goal is to provide a drawing interface on mobile devices that provides the feeling that the user is in full control, while simultaneously providing assistance, in particular, to overcome the inherent fat finger problem. To provide assistance, we take advantage of previous drawings of the same face and then pull the users strokes towards a *consensus* of strokes from previous drawings. We develop a stroke correction strategy with two phases.

(i) Consensus Finding: Using the training drawings available for an image, we create a correction vector field which indicates, for each pixel on the image, the delta toward the nearest *consensus stroke*. This phase is run off-line.

(ii) Interactive Correction: The correction vector field is transmitted to the mobile device along with the image to be traced. When a user draws a stroke, the field is sampled, and

the stroke is moved in real-time in a way that maintains the original style.

Results

After smaller-scale trials, *DrawAFriend* was released on January 8th, 2013. In four days players downloaded the game over 2000 times and created 6373 drawings. In that time, players had already spent 10 full 24 hour days drawing. By July 24th, 2013 there were 17,370 drawings.

Crowdsourced User Study

One key advantage of developing *DrawAFriend* as an online game is the ability to quickly deploy a study of users at scale. To test the effectiveness of the stroke auto-correction, we instrumented the game to enable a simple AB study. Every time a user started a drawing they were unknowingly placed in one of two groups. One group drew the celebrities as before. The other group (unbeknownst to them) drew with the stroke auto-correction on. By comparing these two groups, we can assess the effectiveness of the auto-correction helper.

After approximately one week, 500 players had “contributed” over 1,300 drawings to the user study. Using this data we assess recognition rates (measuring how good the drawings are from the perspective of others). We also analyze undo rates (providing an insight into how artists like their own drawings). Finally we measure average “distance” from the artistic consensus (measuring how carefully artists are drawing). First we investigate *undos* as an indication of how artists like their own drawing. We found our auto-corrector lowers the undo rate for more careful drawers (players that undo a larger percent of their strokes). With auto-correct on, we observe a significant increase in the average distance between the actual uncorrected strokes and the consensus drawings (as measured by the *correction vector field*). This suggests that artists do not need to draw as precisely. Interestingly, statistical analysis reveals that the auto-corrector does not significantly alter recognizability.

Conclusion

We have presented a unique crowdsourcing approach using social game mechanics to grow and use a dataset of drawings. We developed an iPhone game, *DrawAFriend*, to collect drawing data, and then introduced a method to extract stroke-level artistic consensus from a large drawing corpus. The resulting *correction vector field* improves strokes in real-time without new interactions interfaces, while preserving artistic intent. Lastly we evaluated our stroke correction algorithm through crowd sourcing, by observing how it influenced drawings “in the wild” by instrumenting the game itself. In general, we believe that *DrawAFriend* presents an unprecedented platform to perform quantitative drawing analysis at the Internet scale.

Stroke correction is just the tip of the iceberg for applications of the large (and growing) *DrawAFriend* image corpus. We have explored only a small portion of the trove of data we are collecting. Making this data available to the community, we hope to explore both these exciting ideas, and discover as-yet unknown applications of this rich dataset.