My Text in Your Handwriting

Tom S. F. Haines, Oisin Mac Aodha and Gabriel Brostow
{thaines,o.macaodha,G.Brostow}@cs.ucl.ac.uk

27^{th} July 2016

Problem

- To synthesis *user specified text* with a *specific authors handwriting*.
- To a standard that can fool human beings.
- This is guided texture synthesis.

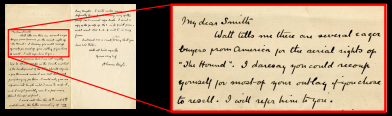

Original letter Tagged paragraph

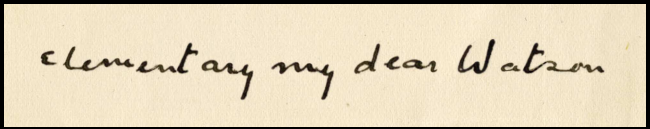

Synthesised Result (Arthur Conan Doyle)

Why?

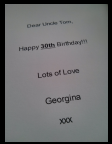

Gift cards/flowers sent without the giver present (e.g. via Amazon). Either use handwriting model of giver or a celebrities.

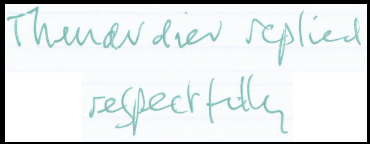

Synthesise handwriting when a disability affects writing. Build model with pre-disability samples.

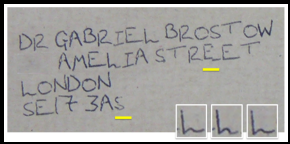

Camouflaging important documents in the mail, such as bank cards.

Artistic purposes, such as comic book lettering. Computer games in particular can contain large quantities of text.

The Problem With Graphics Tablets

- Almost all previous methods use graphics tablets for input.
- Distorts the authors handwriting – like a whiteboard.

Real

Tablet

(Written by an experienced tablet user)

- Has advantage of actual pen path, but disadvantage of no texture.
- Can't use historical writing.
- Does not satisfy use cases – we therefore scan normal writing on normal paper.

Preparation Overview

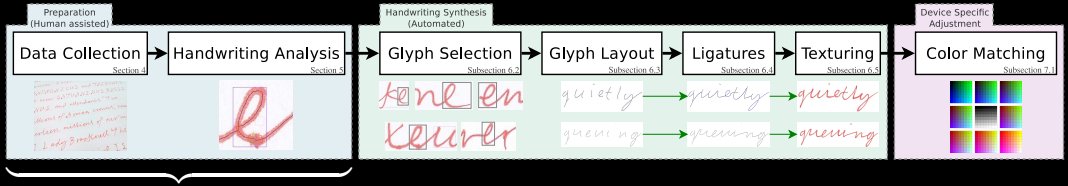

- Input:
  - A willing *author*.

    or
  - Examples of an *author*s handwriting.
- Output:
  - An authors *tagged handwriting*.

Data Collection

- Author unavailable $\rightarrow$ Scan in whatever is available.
- Author available $\rightarrow$ Get them to write out a sample:
  - We can optimise! Maximise coverage for number of words written.
  - Use proper sentences – unfamiliarity changes a persons writing.
  - Select them to minimise extended Scrabble scores.
  - Need a corpus (sentences and statistics) – top 100 books from Project Gutenberg.

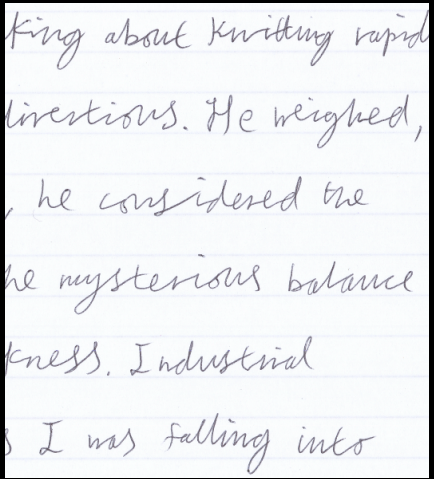

Extended Scrabble Scores

- Scrabble score: $\sim\propto$ number bits required to encode each letter (Shannon entropy).
  - Original statistics from front page of *The New Your Times* – we use statistics from Project Gutenberg corpus.
  - We don't quantise.
- Extended to also include pairwise statistics – how people write a letter is influenced by the letters position relative to others (even if print, and especially start/end of word).

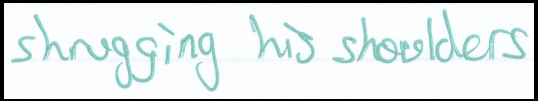

- Extended to consider sentences already selected, so we get variety.

Analysis

- We need to obtain information about scanned handwriting sample:
  - The Rule.
  - Segmentation.
  - Alpha Matting.
  - Spline Fitting.
  - Glyphs / Ligatures.
- We automate as much as reasonable, but allow human intervention.

The Rule

- Rule – line on which the author is writing on.
- On back of page so barely visible in scan (if author available).
- User click-drags to set a homography.
- Not worth automating.

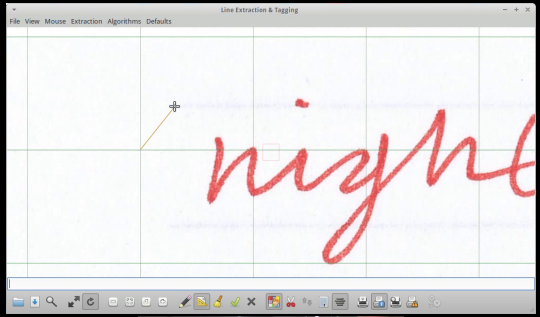

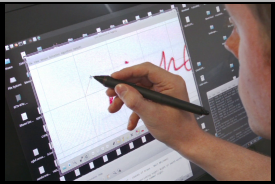

Segmentation

Separate ink line from background:

1. Use mean shift on RGB cube:
   - Largest mode is background colour.
   - Second largest mode is ink colour.
   - Remaining modes are ignored.
2. Perform graph cuts:
   - Unary term from mean shift.
   - Pairwise term from colour difference.
3. Finally, perform a line aware smoothing.
   - Convert mask to signed distance function.
   - Calculate gradient at each pixel.
   - Smooth in gradient direction only.

User can force pixels to be foreground/background.

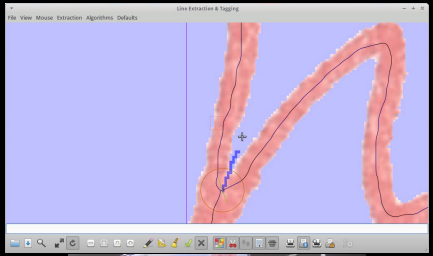

Alpha Matting

- Need to composite output onto arbitrary backgrounds.
- Use segmentation and inpainting to generate a background-only plate.
- Solve matting equation per-pixel to obtain base colour with alpha channel.
- Automatic only.

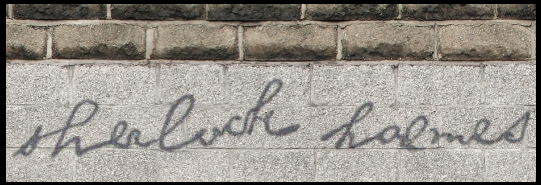

(Arthur Conan Doyle's handwriting)

Spline Fitting

- Thin mask to extract line.
- Assign radius as largest circle that fits within mask at each pixel on the line.
- Assign density as mean unary term from segmentation within radius.
- Automatic only.

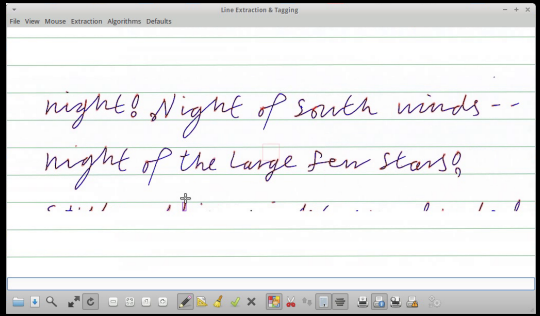

Glyphs / Ligatures

- Attach meta data to the line:
  - *Splits* mark the transitions between glyphs / ligatures.
  - *Links* indicate two separate lines are part of the same glyph (e.g. tittle and stem of “i”)
  - *Labels* indicate which UTF-8 character code a glyph represents.
    (Ligatures are implicitly the lines that attach characters. Start and end of word are also indicated.)
- Automatic system solving constrained handwriting recognition problem.
- Manual editing, as automatic system is not 100% reliable.

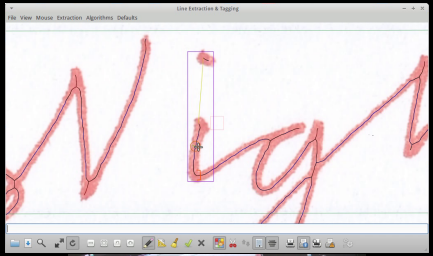

Analysis Summary

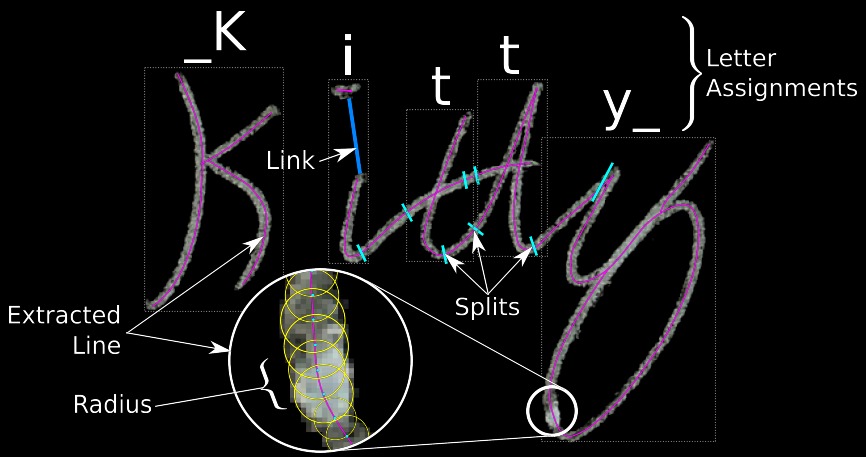

Not shown:

- Ink density
- Matting

Synthesis Overview

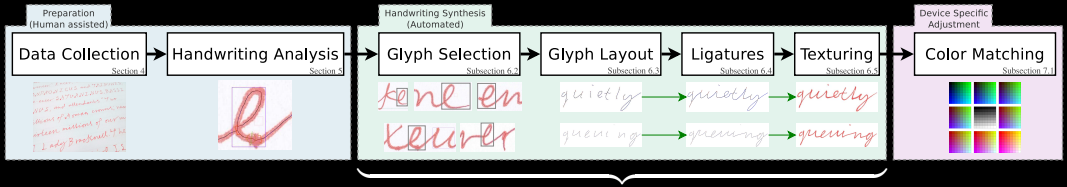

- Input:
  - User provided *text* to generate.
  - An authors *tagged handwriting*.
- Output:
  - A *texture* containing the users *text* written with the *tagged handwriting*.

Core Idea

- *Glyphs* taken directly from the author are positioned on the page.
- The user provides the text, but there is still a choice of glyphs for each position in the word.
- *Spacing* between glyphs – both horizontal and vertical is important to replicating an authors style.
- *Ligatures* need to be generated when appropriate.
- Ultimately, it all has to be rendered to a *texture*.

A cost function is minimised...

Cost Function

We minimise:

$$C(R, t, A) = G_A(g, t) + S_A(g, x) + L_A(g, x, l) + T_A(g, x, l, R).$$

- R – Output texture.
- t – Input text string to synthesis.
- A – Input authors tagged handwriting.
- g – Glyphs to use, one per character in t.
- x – Positions of glyphs, one per glyph.
- l – Set of ligatures to generate between glyphs.

- $G_A(g, t)$ – Match author's glyphs to user's text.
- $S_A(g, x)$ – Match spacing of glyphs to author.
- $L_A(g, x, l)$ – Match ligature use to author.
- $T_A(g, x, l, R)$ – Match output texture to glyphs/ligatures.

$G_A(g, t)$ – Glyph selection

- Constraint that selected glyphs represent the user-requested characters.
- For historical cases there may be missing characters – allow substituting a lower case letter for an upper case letter if necessary.

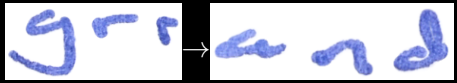

$S_A(g, x)$ – Glyph layout

- Both horizontal and vertical offsets are required between glyphs.
- Original glyph offsets are used where available (ligatures).
- When not available a regression forest estimates them (no ligatures).
- Humans have a feedback mechanism – if you write with your eyes closed you will drift off the rule, but eyes open and you correct for any drift.
- This feedback mechanism is replicated using Kalman smoothing.

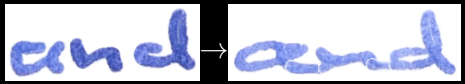

$L_A(g, x, l)$ – Ligatures

- Ligatures are required if both adjacent glyphs have them, but omitted if either is missing one.
- Authors have print handwriting, joined up handwriting, or partially joined up. First two cases are covered by this rule; for third it's a reasonable guess as we don't have enough information to infer exact rules.

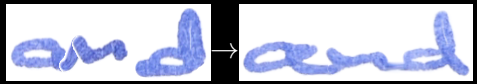

$T_A(g, x, l, R)$ – Texturing

- Encourages nearby glyphs to have similar ink density and radius – sudden changes are unrealistic.
- Hides the seams between glyphs / ligatures (graph cut textures).

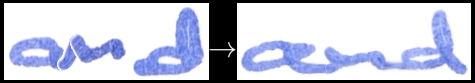

Solving

- Can't solve directly.
- Instead, solve it in four stages, one per equation, in order given.
- Each stage fixes some details using an appropriate representation and approximating other costs with proxies.

Stage	Fixes	Core Technique
Stage 1	Glyphs used (g)	Dynamic programming
Stage 2	Glyph positions (x)	Kalman smoothing
Stage 3	Ligature existence (l)	Heuristic
Stage 4	Output texture (R)	Graph cuts

- Random/regression forests and heuristics are used to estimate proxy costs.

Colour Calibration

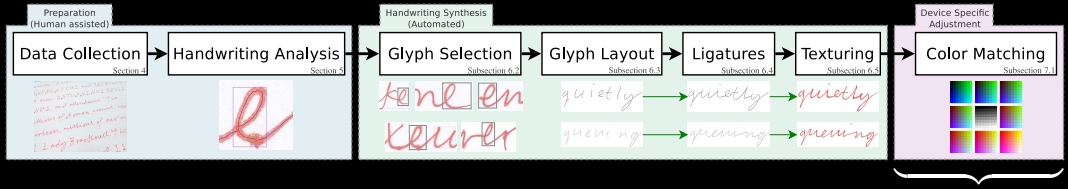

- If printing we need to get as close as possible to real ink.
- Perform a closed loop colour calibration by printing out and then scanning back in a calibration target.
- Uses thin-plate splines.
- Printers lack the dynamic range to cover most inks, but it still helps fool an observer.