

# Imprinto: Enhancing Infrared Inkjet Watermarking for Human and Machine Perception

**Martin Feick**

DFKI and Saarland University,  
Saarland Informatics Campus  
Saarbrücken, Germany  
MIT CSAIL  
Cambridge, Massachusetts, USA  
martin.feick@dfki.de

**Xuxin Tang**

Virginia Tech  
Computer Science Department  
Blacksburg, Virginia, USA  
MIT CSAIL  
Cambridge, Massachusetts, USA  
tangxxwhu@gmail.com

**Raul Garcia-Martin**

University Group for Identification  
Technologies  
Universidad Carlos III de Madrid  
Leganes, Madrid, Spain  
MIT CSAIL  
Cambridge, Massachusetts, USA  
raulgarc@ing.uc3m.es

**Alexandru Luchianov**

MIT CSAIL  
Cambridge, Massachusetts, USA  
lknv@mit.edu

**Roderick Wei Xiao Huang**

MIT CSAIL  
Cambridge, Massachusetts, USA  
rwxhuang@mit.edu

**Chang Xiao**

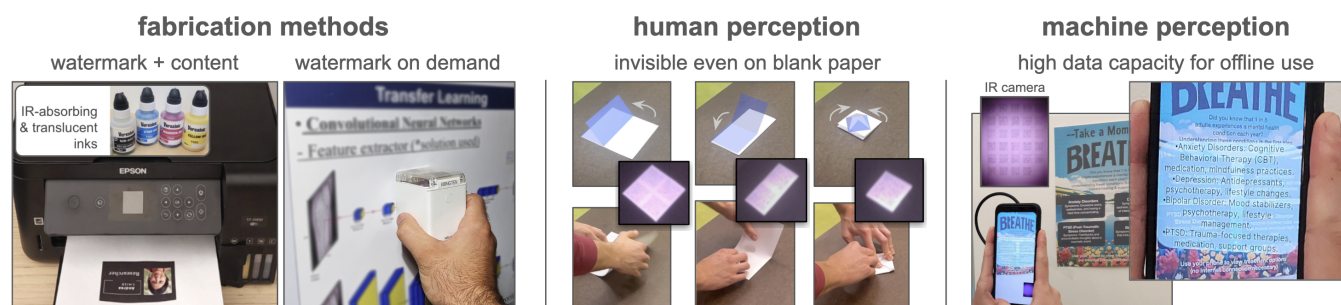
Adobe Research  
San Jose, California, USA  
cxiao@adobe.com

**Alexa Siu**

Adobe Research  
San Jose, California, USA  
asiu@adobe.com

**Mustafa Doga Dogan**

Adobe Research  
Basel, Switzerland  
MIT CSAIL  
Cambridge, Massachusetts, USA  
doga@adobe.com



**Figure 1: IMPRINTO** are real-life documents that enable watermarking through the use of IR-absorbing ink. Users can apply it by (1) making it part of the document's fabrication process using an off-the-shelf inkjet printer or (2) by using a handheld inkjet printer on demand. The embedded content is invisible to users, even on blank paper, while allowing for high data capacity. Our system reliably decodes the IR content and is *further instructed by it* regarding what images to track and what to display.

## Abstract

Hybrid paper interfaces leverage augmented reality to combine the desired tangibility of paper documents with the affordances of interactive digital media. Typically, virtual content can be embedded through direct links (e.g., QR codes); however, this impacts the aesthetics of the paper print and limits the available visual content space. To address this problem, we present IMPRINTO, an

infrared inkjet watermarking technique that allows for invisible content embeddings only by using off-the-shelf IR inks and a camera. IMPRINTO was established through a psychophysical experiment, studying how much IR ink can be used while remaining invisible to users regardless of background color. We demonstrate that we can detect invisible IR content through our machine learning pipeline, and we developed an authoring tool that optimizes the amount of IR ink on the color regions of an input document for machine and human detectability. Finally, we demonstrate several applications, including augmenting paper documents and objects.



This work is licensed under a Creative Commons Attribution 4.0 International License.  
CHI '25, Yokohama, Japan  
© 2025 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-1394-1/25/04  
<https://doi.org/10.1145/3706598.3713286>

## CCS Concepts

• **Human-centered computing** → *Human computer interaction (HCI).*

## Keywords

augmented reality; mixed reality; infrared imaging; watermarking; digital fabrication

### ACM Reference Format:

Martin Feick, Xuxin Tang, Raul Garcia-Martin, Alexandru Luchianov, Roderrick Wei Xiao Huang, Chang Xiao, Alexa Siu, and Mustafa Doga Dogan. 2025. Imprinto: Enhancing Infrared Inkjet Watermarking for Human and Machine Perception. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26–May 01, 2025, Yokohama, Japan. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3706598.3713286>

## 1 Introduction

Despite the growing adoption of digital devices, the use of paper documents remains prevalent in our daily lives [26]. However, paper-based information can often lack the interactivity and accessibility of digital media [31]. In recent years, various paper augmentation techniques have been proposed to bridge this gap, enabling users to enhance the information content of their documents by seamlessly integrating digital media [29, 50, 60]. However, the main disadvantage of these techniques is that the printed markers interfere with printed paper media. As a result, they can affect the aesthetics of the print, breaking the illusion that digital is seamlessly embedded in the physical. This goes as far as distracting consumers from the relevant parts or even degrading overall readability and accessibility. This has prevented techniques from going beyond single links or QR codes hidden at the bottom of the document to seamless paper document augmentations with rich digital content [17].

To address this challenge, we present IMPRINTO<sup>1</sup>, a novel infrared inkjet watermarking technique that allows for invisible paper augmentation. IMPRINTO enables users to interact with printed documents in new ways by embedding hidden high-capacity digital information that can be accessed using a smartphone equipped with an infrared (IR) sensor. The technique works by printing an invisible watermark using an IR-absorbing ink. When the paper is illuminated with IR light, the watermark becomes detectable and can be captured by an IR camera. In contrast to previous work, **our watermarking approach leverages the whole document area, including text, graphics, and white space, to increase capacity and support the embedding of a larger amount of information while remaining invisible to the naked eye.** By utilizing commercially available inkjet printers and inks, we rapidly fabricate these encoded documents without having to include complicated or manual processes such as spraying [76] or screen printing [41].

To inform our system, we first conducted a psychophysical experiment to determine estimates for the maximum amount of IR ink we can embed in paper documents while preventing the IR watermarks from becoming visible to users. Our results show that IR ink can remain invisible to the human eye, regardless of the background color, i.e., even on white paper. Next, we demonstrate that invisible document watermarks can be captured and decoded using our mobile IR sensor and a machine learning (ML) pipeline based on a convolutional neural network (CNN), which binarizes IR watermarks into decodable black-and-white patterns. Our technical evaluation investigates various factors related to the IR ink machine

detectability and data capacity, providing guidelines for content creators and designers who want to use our invisible watermarking technique. Together, our results demonstrate that we can embed invisible content into paper documents that can still be decoded by an off-the-shelf hardware and software stack.

To facilitate the embedding process, we developed a software tool that uses the determined IR ink values to embed digital content invisibly onto paper documents in the form of QR codes. Here, a user simply uploads an image, specifies the number of watermark embeddings, and marks where these should be located. The software analyzes the background colors given the provided embeddings and paper regions, recommending the amount of IR ink that can be used while remaining invisible. The user can then directly fabricate the document using conventional inkjet printing (see Figure 1).

We demonstrate the versatility of IMPRINTO through a series of applications, which highlight a variety of use cases covering education, arts and design, desktop publishing, robotics, entertainment, sustainability and security. Unlike other paper-based augmentation techniques, **IMPRINTO is invisible, works with new and existing documents, and does not require any embedded electronics or special paper.**

To summarize, the contributions we make are as follows:

- IMPRINTO, an inkjet printing-based watermarking method that allows the utilization of the whole sheet (both blank and color-printed areas) for invisible data embedding onto new and existing physical media.
- A methodology on how to evaluate human perception of IR inkjet ink, and its optimization for invisibility, data capacity and machine detection.
- A universal hardware detection module compatible with USB-C devices.
- An open-source<sup>2</sup> machine learning pipeline for binarizing the captured IR watermarks for robust detection and data capacity under different visible background content and IR ink levels.
- A set of use cases that demonstrate IMPRINTO's application space, even beyond paper-based documents.

## 2 Related Work

Combining the digital and the physical worlds has been one of the main uses cases of augmented reality [11, 16]. Below we explain how the two worlds interact with each other, how paper specifically is used for these interactions, and how tags are used to specify interactions in previous works.

### 2.1 Hybrid Paper-Digital Interfaces

The benefits of hybrid paper-digital interfaces have been extensively explored in prior work [1, 50, 60, 66]. Han et al. conducted a systematic literature review and defined hybrid paper-digital interfaces as *"any interface embedding digital or electronic functionality in physical paper to enable its use as an input or output device"* [29].

To enable input and enhance digital paper applications, researchers proposed integrating electronics into paper sheets, such as sensors [27, 34], actuators [58], or illuminated displays [44, 45].

<sup>1</sup>The system name IMPRINTO is a combination of *imprint* and *pronto* (quickly), indicating that existing documents can be quickly watermarked for additional AR content or metadata embedding.

<sup>2</sup><https://imprinto.github.io>

Buechley et al. further proposed *paper computing* [4] as a concept to allow users to create functional computational artifacts on painted paper substrates using a combination of these paper-embedded components as well as microcontrollers. On the other hand, leveraging AR is a common technique used to add a digital information layer but avoid integrating electronics into paper. For instance, *Dually Noted* [59] and *HoloDoc* [50] demonstrate the utility of augmenting paper for improving productivity tasks for phone- or headset-based AR. *PapARVis* [6] and *Paper Trail* [60] explore authoring workflows that support educational and data visualization use cases. *PaperLens* [67] tracks sheets in 3D space above a top-projected tabletop [68] to visualize and interact with 3D information spaces. *SpARKlingPaper* [21] augments paper sheets from below using a tablet's screen placed underneath. The paper augmentation in these projects are enabled by systems such as dedicated standalone apps that were tailored for specific use cases. Most of these leverage image tracking libraries such as *Vuforia*<sup>3</sup>, *ARKit*<sup>4</sup>, and *ARCore*<sup>5</sup>. These utilize visual features (*SURF*, *SIFT*) or fiducial markers (*ArUco* [24], QR codes<sup>6</sup>, *ARToolKit* [39]), which is based on previous research in computer vision (CV) [70, 71]. While visual feature based tracking does not require the addition of new visual elements such as QR codes, they do not carry high capacity to be able to represent a **unique and direct anchor to digital experiences**. In contrast, *AniCode* [73] uses QR codes to store entire animation sequences contextual to an object, without having to access digital content from an external service. A remaining limitation is that the visible QR codes become obtrusive in the user's view as longer sequences are embedded.

In this work, we investigate invisible IR watermarking as the fabrication method to enable hybrid documents to be portable containers of *both* physical (visual) and digital (invisible) content. Our aim is to preserve documents' defining attribute of portability by making its digital link **an intrinsic part of the document**.

## 2.2 Linking Digital Information to Physical

To connect paper documents to digital content, machine-readable printed codes, such as barcodes, have been actively used as a compact and cheaper way to embed digital information [10]. 2D barcodes, such as QR codes, store information in the form of contrasting bits, but affect the original design as they occupy space and impact the visual look of the final artifact [25, 65], particularly those that store large amounts of data [75]. Thus, researchers explored making more unobtrusive camera-detectable tags to create a more seamless, discreet user experience [12, 14]. For this goal, works investigated subtle pixel manipulations to encode data *unobtrusively* [23, 69, 77, 79]. For instance, *Chartem* [23] encodes subtle pixels into chart background, however, it only supports digital versions of charts, i.e., the data cannot be robustly decoded once they are printed. Methods that support printed documents can only leverage colored pixels and are thus limited in the amount of stored information. For instance, *StegaStamp* [69] can encode up to 7 characters into a single image, and *FontCode* [77] encodes 1.77 bits per character (e.g., a 1,000-character paper abstract would store 221

encoded characters). Furthermore, they cannot store information in non-printed white space, which is a typical part of documents.

IMPRINTO leverages the **whole document area, including text, graphics, and white space, to increase capacity and support the embedding of digital (invisible) content regardless of the physical (visual) content layout**.

## 2.3 Aesthetically Integrated Visual Markers

Another relevant approach is to make the markers human-designable so that they seamlessly blend into the aesthetics of the content. This relies on visual features that follow a predefined set of rules, allowing the system to decode the integrated marker. For example, *d-touch* [8] uses the relationship of dark and light regions in the image for binary encoding. While the initial technique was limited to smaller areas and black and white, Preston et al. [56] extend this approach to impressive scale and aesthetics by adding area order codes and visual checksums. Benford et al. [3] went beyond paper-based interfaces by embedding computer-readable codes in the form of decorations to ceramic bowls, fabric souvenirs, and an acoustic guitar, demonstrating the wide application space.

Studies showed that designers can incorporate these high-capacity codings in their visual designs [8, 35]. However, it may require additional mental effort, be prone to errors, and limit artistic freedom. They are further challenged by the necessity to adapt the visual design to specific rules and remain limited to areas displaying content. Moreover, visual markers can only be used during the creation process [36], whereas IMPRINTO *can be used as part of the fabrication process or added on demand*. Nevertheless, this impressive line of research demonstrates how high-capacity markers can seamlessly blend into artistic images, inspiring us to look into alternative approaches that allow invisible data embeddings while preserving the aesthetics of new and existing content. Ultimately, we hope the techniques could complement one another to allow for even higher data capacities.

## 2.4 Infrared-Based Tags

IR based approaches such as patterns projected by external sources [37, 47], materials that pass [18], reflect [2, 41], absorb [63, 76], or fluoresce [15] IR light to embed patterns have emerged. For instance, *MiniStudio* [41] tracks objects by adding IR-reflective stickers with fiducial markers using screen printing. *InfraredTags* [18] covers fiducial markers with an opaque material that passes IR light to make them imperceptible.

As for IR-absorbing materials, *Spyn* [63] adds such inks into knitted artifacts to embed hidden information. *HideOut* [76] creates hidden tracking markers from IR-absorbing ink to project digital imagery on physical paper; however, a spray gun is needed to evenly coat the paper surface. Wang et al. [72] showed how different IR absorption characteristics of CMY and black ink can be used to subtly embed fiducial markers; however, it is limited to dark areas<sup>7</sup>. Similar to this work, we leverage the convenience and accessibility of commodity inkjet printers, commonly used in ubiquitous computing research [7, 40], which allow us to print hybrid interfaces rapidly using the encoded data, but IMPRINTO *also makes use of the empty space in the document*.

<sup>3</sup><https://developer.vuforia.com/>

<sup>4</sup><https://developer.apple.com/augmented-reality/arkit/>

<sup>5</sup><https://developers.google.com/ar/>

<sup>6</sup><https://www.qrcode.com/en/>

<sup>7</sup><https://frauzfall.de/en/2020/print-invisible-ir-markers-using-a-standard-printer/>






Figure 2: Spectral analysis of the inkjet inks.

To summarize, IMPRINTO has three main benefits that allow it to rapidly create portable hybrid document experiences: (1) We utilize the whole document area, including text, graphics, and white space to increase the freedom of placement and amount of data that can be invisibly stored in new and exciting content. (2) We can rapidly create these encoded documents without having to use complicated or manual processes such as spraying or screen printing, thanks to commercially available inkjet printers and inks. (3) We build on previous demonstrations [17] and establish a scientific basis on how to use perception research to truly hide digital IR content for different classes of visibly printed content, while optimizing for machine detection and information capacity using ML.

### 3 Infrared-Absorbing and -Translucent Inkjet Inks

One key component in our system is the set of inks used in our printed samples. For our implementation, we used an off-the-shelf desktop inkjet printer (*Epson ST-200*), matte paper from *Canon*<sup>8</sup>, and four types of inkjet-compatible inks (C, M, Y, IR) by *VersaCheck*<sup>9</sup>. The desktop printer has four refillable tanks, which we leveraged for visible and invisible content printing. The cyan, magenta, yellow (CMY) channels of the printer were filled with the infrared-translucent C, M, Y inks, and the *K* (i.e., conventionally black) channel was filled with the IR ink (NIR, 78 ml bottle). This ink passes light at visible wavelengths, and thus appears mostly transparent to the naked eye, but has high absorbance intensity in the near-infrared (NIR) range, with the peak absorbance wavelength at 850 nm. Thus, it is captured as black under an infrared camera lens. Alternately, users can leverage the flexibility of handheld printers [55] such as the *Kongten MBrush*<sup>10</sup> to embed watermarks on demand, only requiring IR ink. However, we used the desktop printer for our investigations of human and machine perception as it allowed us to produce very consistent and reproducible results.

As the ink loaded into the "black" channel is non-black (for human sight), we instead generate the black visible content by

mixing the C, M, and Y inks, as their combination results in black color. We achieve this elimination of the black channel requirement by converting the original RGB color space of the visual part of the document to CMY (i.e., as opposed to CMYK) by setting the *K* channel to zero through Gray Component Replacement (GCR) [38] via *Adobe Photoshop*. We note that the resulting black color may slightly visually differ from off-the-shelf black ink. The visible color content of the document is printed this way initially, then the sheet is fed back to overlay the invisible content with the NIR ink. The watermark's content is then printed using the *K* channel only.

#### 3.1 Absorbance Analysis of Inks

We conducted a spectral analysis on the four types of inks we use in our system (infrared-translucent C, M, Y, and infrared-absorbing IR) with the help of a spectrophotometer, the *Perkin Elmer Lambda 1050 UV/VIS/NIR*<sup>11</sup>. In addition, we also measured the properties of *conventional* CMY inks, so we can compare it them to our *infrared-translucent* CMY inks. Our results are shown in Figure 2. The x-axis represents wavelength of light measured in nanometers (nm), and its range 380-1100 nm was selected to illustrate the behaviour under visible and NIR light conditions. The y-axis is  $\log(1/R)$ , which indicates the absorbance rate calculated via the reflection (*R*) measured by the device.

Within the human visibility range (from 380 nm to 750 nm), both regular and infrared-translucent CMY inks are more discernible to the human eye compared to the IR ink, which registers close to zero visibility and is therefore harder for humans to perceive. However, beyond 750 nm, the absorbance of conventional CMY inks drop to around 0.1, while those for infrared-translucent CMY inks drop closer to 0. At around 800 nm, the absorbance of the IR ink surpasses that of the three infrared-translucent CMY inks, as shown in the callout of Figure 2. This means that starting at this wavelength, the IR ink would look "dark" (i.e., absorbing) to an IR camera in comparison to translucent CMY inks. Notably, at 850 nm, the difference between the IR ink and infrared-translucent CMY inks is the highest. The larger this difference, the easier it

<sup>8</sup><https://www.amazon.com/Canon-7981A004-Photo-Paper-Sheets/dp/B0000721Z3/>

<sup>9</sup><https://www.versacheck.com/>

<sup>10</sup><https://www.amazon.de/-/en/EVEBOT-Kongten-Mbrush-Printer-Android/dp/B0C84P61WH?th=1>

<sup>11</sup><https://www.perkinelmer.com/product/lambda-1050-2d-base-inst-no-sw-16020055>






Figure 3: User study samples.

is to distinguish the watermark from the the CMY content under an IR camera. Thus, at this wavelength, one could more reliably decode the data printed with IR ink, as the CMY content would look most dispersed. Thus, we use IR ink in combination with these IR-translucent CMY inks. In comparison, conventional CMY inks are even "darker" than IR ink at this wavelength, meaning they would block the IR content from being detected if they were used for printing the documents instead.

Consequently, IR cameras should be calibrated to image at 850 nm with the background printed using our infrared-translucent CMY inks to enhance the detectability of the IR ink.

#### 4 Experiment: Detectability of IR Ink

One challenge with IR ink is that, even though it allows for invisible pattern embedding, it may become visible to the naked eye when its concentration is too high. This results from a difference in intensity, that is, a contrast between the IR-inked parts and the rest of the paper region. Consequently, it can potentially affect the readability and aesthetics of the print. Our central interest lies in using the highest concentration of IR ink to improve the detection rate while remaining invisible to human eyes. The experiment was conducted indoors under normal office lighting conditions (i.e., well-lit, about 300 lux and a color temperature of 3500K–5500K).

To investigate this, we conducted a psychophysical detection threshold experiment on the detectability of printed IR ink at different background colors. We included colors with varying RGB values, darkness ( $K$ ) [49] and of different luminescence [61] to cover a large color spectrum and to understand how these color properties affect the detectability. This decision was motivated through multiple feedback loops and informal pilot testing, suggesting, e.g., that darker colors mask the IR ink, whereas IR ink on brighter backgrounds can be more easily noticed by humans. To improve the quality and robustness of the data collection, we ran the experiment in the lab under controlled lighting conditions.

We formulated the following two hypotheses in our experiment:

**(H.1)** Inkjet printed IR ink can remain invisible to human eyes, regardless of background color.


**(H.2)** The maximum density of IR ink that remains invisible depends on the properties of the background color.

#### 4.1 Method

To determine estimates for detection thresholds, we used the psychophysical methods of limits [42], exposing participants to ascending stimuli, that is, increasing the concentration of IR ink, until they noticed it. Participants were presented with two color bars, one with IR ink and the other without IR ink, and we asked them to compare both color bars and report if they noticed a difference between them using a 2-alternative-forced-choice (2AFC) method. As depicted in Figure 3, we used a gradient pattern of IR ink, starting with no IR ink 0% (left) to 200% (right) density. By using 0% as the starting stimuli, we ensure that participants start with a level of IR ink that cannot be noticed [48]. The higher IR ink densities of 101% to 200% require two successive printing procedures. Given the novelty of our approach, we determined a reasonable range of IR ink densities through pilot testing.

**Task.** Participants moved a thick white paper card (depicted in Figure 5) from the left (starting stimuli) to the right. We asked participants to move the card slowly and evenly, stopping and reporting immediately when they detected the IR ink. The experimenter recorded the position at which they stopped using the ruler printed at the top of the page (see Figure 3), which was not visible to the participants. Consequently, participants continued with the next row until they completed all rows. Only one row was visible to the participants at a time, and everything else on the page was covered. Participants were informed about the procedure, and we explicitly showed them IR patterns multiple times during the warm-up phase to ensure that they were able to detect it. We asked them to report the IR ink as soon as they noticed it, thus targeting the most conservative case. Participants took breaks whenever needed.

**4.1.1 Background Color Selection & Preparation.** Our pilot studies suggested that the detectability of IR ink likely depends on the background color. Therefore, we included colors with varying RGB values, darkness ( $K$ ) and of different luminescence to obtain a diverse set of colors. To cover a wide range of colors with different levels of darkness and luminescence, we first divided the darkness (calculated as an analog  $K$  value from RGB) into four equal 25% intervals, ranging from 0% to 100%. We then calculated the range of luminescence covered within each interval. As the darkness of a color increases, the luminescence range narrows, meaning the



**Figure 4: IR ink DTs for each color tested in our experiment, reaching from a mean DT of 81% all the way to 200% IR ink density. The three classes are provided based on participants’ ability to detect the IR ink.**


colors become darker. Therefore, we allocated four colors to each  $\sim 25\%$  luminescence interval, starting with the lightest colors in Group 1, where darkness is between 0–25%. Group 1 and Group 5 include the two extremes, no background color and black. Finally, we added four additional color hues to cover a large spectrum of real-world use cases and scenarios. A detailed list of the colors tested in our experiment is provided in Table 1 in Appendix A.

Finally, we printed the background colors with the IR ink gradient as described in section 3. The size of the color bars was chosen to give participants a large enough area to visually inspect the color. As a result, we randomly split the 45 samples into groups of eight, resulting in six pages depicted in Figure 3a.

**4.1.2 Design.** We used a design within subjects in which each participant successively inspected each background color bar on all six pages. We fully counterbalanced the order of the six pages using a Latin square ( $n = 6$ ).

**4.1.3 Participants.** We recruited 18 right-handed participants (six females, twelve males), aged 18–31 (mean: 25.05; SD: 3.05) from the general public and the local university. Participants had a range of different educational and professional backgrounds, including computer science, education, cybersecurity, mechanical and electrical engineering, data science, and artificial intelligence. All participants reported normal/corrected-to-normal vision and did not report any known health issues which might impair their perception. Participants received a \$10 Amazon voucher for taking part in the experiment. The study was approved by the MIT’s Ethics Board.

**4.1.4 Experimental Protocol.** Following a general introduction to the experiment, participants completed a demographic questionnaire. We then explained the task by showing them a set of prepared samples with background colors different from those used in the study. Next, we demonstrate how participants should explore the samples, i.e., moving the card from left to right with a slow and consistent speed. They were instructed to sit comfortably with an approximate reading distance of 30–40 cm which mimics a typical real-world reading distance. We decided against fixating participants’ heads, e.g., with a chin rest, to collect realistic real-world data. The experimenter monitored that they stayed within a reasonable distance, i.e., not going too close to the sample. Once the



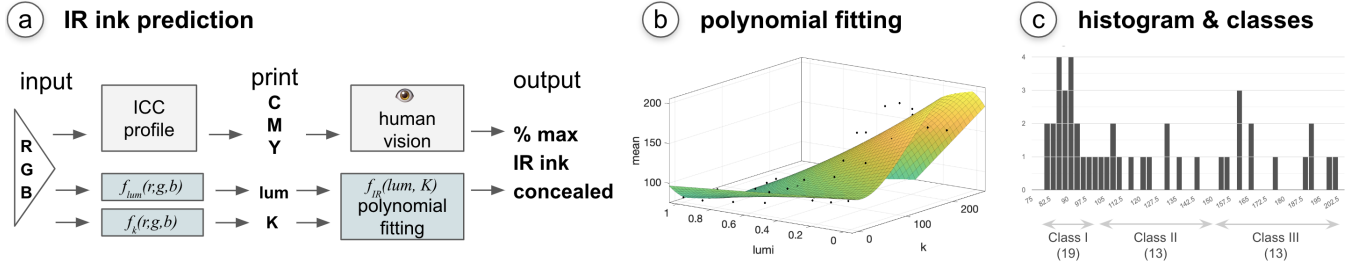
**Figure 5: Study procedure. The participant slides the card steadily to the right to reveal the color bar. When participants noticed a color difference between the top (a) and the bottom (b) of the color bar, they were told to report it.**

participant indicated that there was a color difference, the experimenter obtained the corresponding IR ink value, which remained invisible to the participant. If a participant reached the end of the color bar and did not report a difference, the maximum IR ink value of 200% was used as a response. Participants were given unlimited time to practice to ensure that they understood everything. They were encouraged to ask questions in case something was unclear. In total, the experiment took 35–45 minutes.

**4.1.5 Data Collection & Analysis.** We collected data from three sources: a pre-study demographic questionnaire; observations and field notes, and participants’ responses to the 2AFC method. For each of the 45 background colors, we received 18 responses, resulting in a total of 810 data points. First, we performed outlier removal using the Box-plot method. Next, we statistically analyzed our data after verifying the parametric test assumptions at  $\alpha = .05$  using an RM ANOVA. In the presence of a main effect, we performed post hoc pairwise comparison t-tests adjusted using the Bonferroni-Holm  $p$ -value adjustments.

## 4.2 Results: Detectability of IR Ink

**4.2.1 Printed IR Ink Can Remain Invisible to the Human Eye.** We compute an overall DT for each background color by averaging



**Figure 6: IR ink computation.** (a) We attempt to mathematical model the limits of human perception to conceal IR ink. (b) We use polynomial fitting to estimate the IR capacity based on  $lum$  and  $K$ . (c) The class limits were determined on the histogram of IR ink percentage.

individual DTs of participants. The results can be seen in Figure 4, showing that for each background color, we obtained a DT greater than zero, reaching from DTs of 81% (no background color) all the way to 200% (black background color). This confirms H.1 showing that the inkjet-printed IR content can remain invisible to the human eye. We also observed that the darker colors especially seem to suffer from high variance due to a wide range of participant responses. This may result from two reasons: (1) high-density IR ink has a “greenish/grayish” appearance when printed. Background colors with the highest variance also show this color tone (see Figure 4) and, therefore, may be mistakenly perceived as IR ink. (2) we examined the high-variance samples after the experiment, and even though we fabricated high-quality prints, we noticed slight impurities. Given the nature of the task, i.e., we asked participants to fully focus on the sample and report a color difference as soon as they noticed it, we believe this might have pushed people to pick up on this, leading to even more conservative IR ink estimates.

**4.2.2 Maximum IR Ink Density Depends on Background Color.** Figure 4 suggests that the invisible IR ink density depends on the background color. This is confirmed by our statistical analysis, which reveals the main effect of the background color on DT ( $F(44) = 27.369$ ,  $p < .001$ ,  $\eta_p^2 = .631$ ). 481/990 post-hoc pairwise comparisons were statistically significant after  $p$ -adjustments, demonstrating that background color has a significant effect on the DTs in terms of the invisible IR ink density. Thus, we can also confirm H.2, that the maximum IR ink density depends on the background color.

### 4.3 Predict Invisible IR Ink Density

To facilitate the design process of watermark embeddings using invisible IR ink, we developed a system that estimates IR ink density for a given RGB color input (see Figure 6). First, we performed two correlation analyses for  $K$  and  $lum$  for our 45 background colors on the amount of invisible IR ink to understand their influence. We found that  $K$  shows a strong negative correlation ( $\rho(43) = -.678$ ,  $p = .012$ ) and  $lum$  a strong positive correlation ( $\rho(43) = .775$ ,  $p < .001$ ). This means that increasing  $K$  and decreasing  $lum$  result in higher IR ink densities and vice versa. To give designers and creators a straightforward way to estimate invisible IR ink for any given RGB color input, we fitted a polynomial with 3 degrees of freedom,  $x = K$ ,  $y = lum$  and  $z = \text{mean IR ink}$ , using MATLAB’s polynomial fitting function. The resulting linear model ( $RMSE = 14.48$ ,  $R^2 = 0.8477$ )

that can be seen in Equation 1 outputs an IR ink estimate based on the parameters  $K$  and  $lum$  given an RGB color input.

$$f(K, lum) = 113.7 - 0.1384 * K - 25.43 * lum + 0.0081 * K^2 - 0.5103 * lum^2 + (0 * K^3) - 0.0013 * lum^3 \quad (1)$$

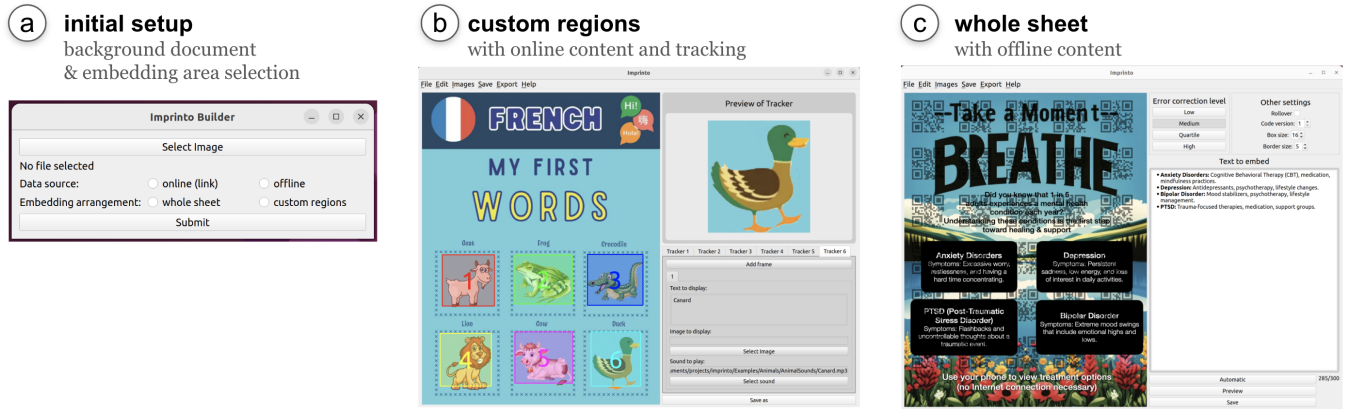
Despite our efforts to include a variety of different colors, the amount of data that can reasonably be collected in psychophysical lab experiments without causing too much fatigue is limited. To improve the reliability of our method for colors that we did not test in our experiment, we first defined three classes based on the histogram of the mean IR ink. The limits of these classes were selected to represent the color distribution based on visual inspection of Figure 6c. Here, we notice that most tested colors are accumulated on the left side, i.e., representing lighter colors. The first class we define captures the accumulation of the lighter colors. In the histogram, the other two classes have the same number of colors to make the binning uniform.

However, designers are free to modify this based on their use case, as our goal here is to provide an example for our applications using our methodology. Class 1 (81% to 101%) contains the colors that allow the least density of IR ink. For all colors that fall into this category, following the polynomial function (1), the system recommends the very conservative value of 81% IR ink (white paper). Similarly, for Class 2 ( $\geq 102\%$  to 154%) it recommends 102%, and 155% for Class 3 ( $\geq 155\%$  to 200%). Actual undetectability thresholds are most likely much higher. Still, we aimed to provide a method that allows designers to include any desired color while ensuring that users cannot notice the IR ink.

### 4.4 Summary

Our results show that the IR ink can remain *invisible to humans regardless of background color* (H.1). Furthermore, we found that the maximum IR ink density depends on the properties, mainly the color  $lum$  and  $K$  of the background color (H.2). We developed a system taking RGB as an input, outputting a conservative lower bound for the invisible IR ink density given a background color. To this end, we provide evidence that IR ink can, in fact, remain invisible to humans; however, it is unclear whether inkjet prints with established IR ink densities can be detected using IR sensors. Therefore, the next section describes how we embedded, captured, and decoded our IR ink watermarks.





**Figure 7: Our software tool for embedding IMPRINTO watermarks. (a) The initial setup window is used to select between online/offline and whole/custom region modes. (b) Online mode allows users to designate custom regions for tracking and embeds invisible QR codes that instruct the app how to access the AR content from the internet. (c) Offline mode optimizes fitting to accommodate as much data on the page as possible.**

## 5 Imprinto Watermarks

In this section, we provide a comprehensive overview of the watermark embedding and detection process. The key steps involved are capturing the image, binarizing the image and extracting the QR code from the binarized image. Although extracting the QR code from an already binarized image can be done using commercially available software, the remaining necessary steps required the development of specialized modules.

### 5.1 Embedding the Watermark

We developed a software tool to embed markers based on the system and the invisible IR ink values determined in our experiment in Section 4. To use our tool shown in Figure 7, users first select the document for which they wish to create embeddings. Then, they choose whether to utilize the whole sheet for embedding or only designated areas. In the latter, the user marks on the document where they want to embed content, and the system interactively generates layouts. Finally, users decide whether to rely on internet information (online) or create a fully self-contained (offline) document. The tool then analyzes the RGB values of the background color of the specified image region. It determines the conservative lower bound of the invisible IR ink density using the underlying system described in Section 4.

*Online mode.* In the online mode, we take advantage of the ability to use the internet by hyperlinking multimedia content using QR codes. In this menu (Figure 7b), the user can select the paper area that they wish to track using AR. An image score is computed for each selection using Google *arcoring* tool<sup>12</sup> to ensure proper tracking and content overlay in AR. In case the image is too difficult to track (i.e., the score is lower than 75), we inform the user and prompt them to choose a different area. Our tool allows users to associate each AR-tracked image with online multimedia content, such as text, image, and/or audio. For each tracked image, the user can link a series ("frames") of these types of contents that

users can loop through in AR. The underlying technology can be further generalized and used in order to incorporate complex multimedia content (such as videos or 3D models) or engineer further interactions between the objects.


*Offline mode.* In the offline mode (Figure 7c), we are currently limited to text-only content due to multimedia content requiring significantly more space. However, this capacity that offline embedding offers can be used to embed hyperlinks or XML media, such as contact cards, or stylized and interactive text content via markup language tags of the user's choice. Even though we offer an automatic analysis of the document and calculate a possible way of embedding the information, we further allow the user to fine-tune these settings (i.e., size, error correction level, and version of the QR codes) according to their specific requirements.

### 5.2 Capturing the Watermark

IMPRINTO media can be captured using devices with NIR image sensors. There are two main methods by which users can do this. First, more conventionally, users can point their devices at the physical media to extract IMPRINTO content. This is particularly suitable for settings where the users are aware that the content is already embedded (e.g., an instructor informing students about IMPRINTO verbally before handing out worksheets, or through a small visual cue on them; a book or magazine mentioning the AR functionality only on its cover—so not each page needs to carry a notice). Second, there are other types of devices that have constant scanning functionality [20, 33]. Emerging devices, such as smart glasses (see *Meta Orion*<sup>13</sup>), feature "always-on" IR scanning functionality to map the environment that we envision can be repurposed to detect human-invisible IR markers (see Figure 13). This, in turn, would enable the always-on and seamless capture of these markers without much physical effort. However, accessibility to the technology is not yet widely available, discussed in Section 8.4.

<sup>12</sup><https://developers.google.com/ar/develop/augmented-images/arcoring>

<sup>13</sup><https://about.meta.com/realitylabs/orion>



**Figure 8: Our mobile imaging attachment comes with a 5MP sensor and two controllable NIR LEDs working at 850 nm. (a) Module directly plugged into the mobile device. (b) Module coupled and lined with the back smartphone camera through a USB extender. (c) Front view of the module with the NIR LEDs powered up on a different device. (d) 3D exploded view, including the sensor, the LEDs mounted in the PCB, and the aluminum mechanical and electromagnetic enclosure.**

To support this emerging technology, we developed a *universal* IMPRINTO detection and reader module for mobile devices. The module was built to decode document watermarks based on the inkjet ink absorbance characteristics (Section 3.1) by featuring a 5MP CMOS NIR camera and two controllable NIR LEDs, both operating at a wavelength of 850 nm. The sensor module integrates a bandpass filter centered at the same wavelength. We designed and developed the PCB and the aluminum enclosure for mechanical and electromagnetic protection produced by Megacal<sup>14</sup>, with a size of 59.5 × 33.0 × 11.6 mm. Figure 8 shows the components of the module. Our proprietary module, *RGMVision InfraRedCAM 1*<sup>15</sup>, ensures universal mobile connectivity through USB 2.0 or 3.0 and its USB-C connector. As opposed to related works, this allows our device to offer universal compatibility, converting any mobile device, smartphone, laptop, or PC into an IMPRINTO reader via its USB-C support. Additionally, the two diffuse NIR LEDs (80°), controlled by two PWM signals, provide 24 different adjustable illumination levels for indoor applications.

We developed an *Android* application (built for API level 26 and targeted for API level 30), based on *Unity3D*, interfaces with the module through the *Unity3D* plugin *UVC4UnityAndroid v2.0*<sup>16</sup>.

### 5.3 Decoding the Watermark

We developed an ML pipeline to binarize the captured IR watermarks for robust detection, which we will make open-source. Because the IR camera captures the watermarks in a faint tone, off-the-shelf QR code readers are unable to recognize IMPRINTO codes. Thus, we need an intermediary step where the contrast of the IR watermarks is increased via binarization, where the codes are converted into black-and-white patterns similar to conventional QR codes. Previous works [17] have used traditional image processing

filters for this step, whereas we use an ML-focused approach to increase robustness (Section 6). We next explain the individual steps that were necessary for building our pipeline, as well as how to use it in real time.

**5.3.1 Training the ML model.** The training involved multiple steps, as shown in Figure 9.

**Data capture.** To ensure our ML model generalized to various conditions, we printed and captured multiple conditions with our setup. We intended to print different types of barcodes so the trained model can excel at binarization as an image task, regardless of the barcode type or content. Thus, we printed samples of three types of barcodes: a QR code with low ECC, a QR code with high ECC, and a Data Matrix code. Per sheet, we printed only one type of barcodes.

On each US letter-size sheet, we had a total of 25 codes that were created by repositioning and rescaling one of these initial barcodes in a grid-like pattern (see Figure 9). Varying the size of the code creates a more robust training set so that the model can recognize codes that appear smaller from a distance, at different angles, or with lower resolution. We varied the the module (bit) size from 1 mm to 2 mm, in increments of 0.25 mm from left to right to allow our model to generalize across different resolutions.

Each sheet was printed three times, once for each of the three possible IR ink quantity classes, and a matching visual for the CMY color content (darker background, lighter background, or blank), resulting in a total of nine sheets. Regardless of the ink amounts and combinations, the model is expected to learn to detect and binarize the codes. We then photographed each sheet using the IR camera from an identical distance of 25.5 cm in order to capture the letter-size sheet (see Section 6.1).

**Ground truth labeling.** Using *Photoshop*, we binarized the images in order to create a black-and-white image representing the ground truth. This process involved manually adjusting each code using various transformations (e.g., perspective warp) to match the pose of the code on the photographed sheet.

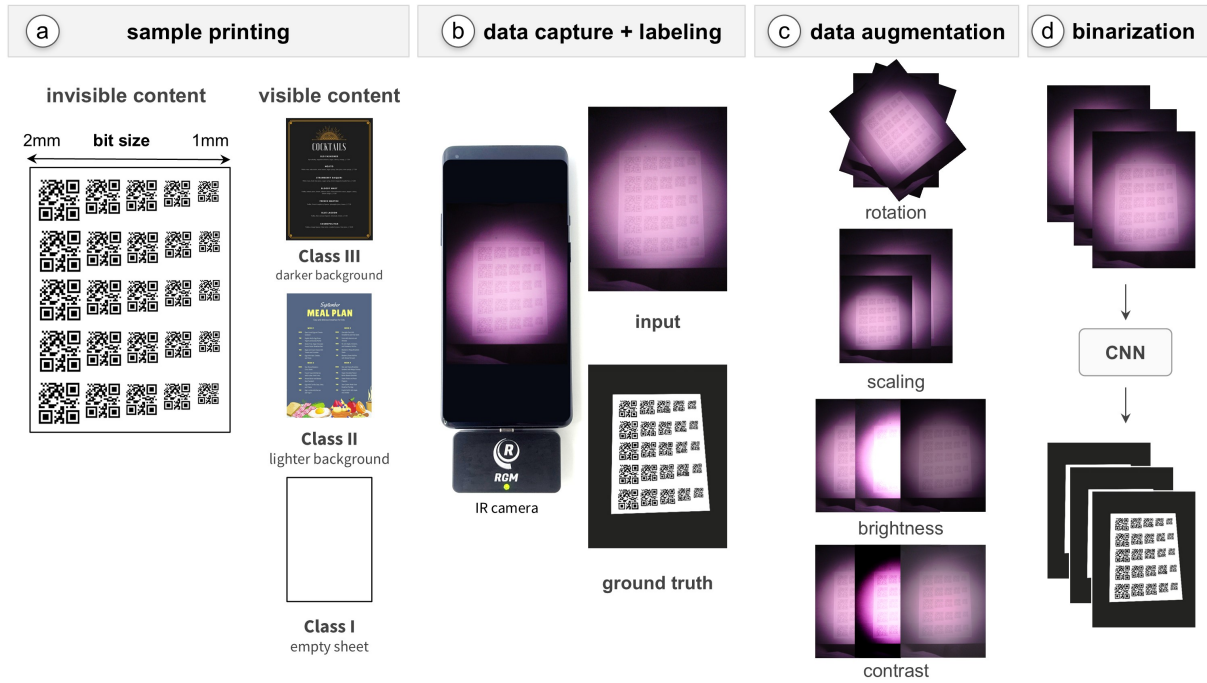
**Data augmentation.** Using *MATLAB*, we augmented our dataset by 500 images for each original image by performing a combination of geometric and color space transformations, including randomized changes in brightness, contrast, hue, and saturation. Figure 9c exemplifies the samples we generate in order to simulate different conditions that might occur due to varying physical conditions. Each image was rescaled to 450x600 pixels for efficient ML training. In total, our resulting dataset had 4,500 images. Such data augmentation in ML also helps us enhance model robustness and prevents overfitting by enlarging the dataset.

**Training.** For the ML training, we employed a convolutional neural network (CNN) similar to [13] but used the *U-Net* [62] architecture with a pre-trained *ResNet-34* [30] as its backbone for black-and-white segmentation (i.e., binarization). Using *PyTorch* [54] and *fast.ai* [32], we utilized the *Adam* optimizer [43], and used cross-entropy as loss function. We trained the model with a batch size of 8 for 30 epochs and for about 20 minutes each, until training loss settled at 0.05. Because this is not a classification task but rather segmentation (binarization), we do not compute an accuracy percentage to evaluate the performance for the sake of the eventual

<sup>14</sup><https://megacal.es/>

<sup>15</sup><https://www.rgmvision.com/product/rgmvision-infraredcam-1/>

<sup>16</sup><https://github.com/saki4510t/UVC4UnityAndroid>



**Figure 9: ML training. (a) Sample printing. (b) Data capture with our smartphone module and labeling processes. (c) Data augmentation. (d) CNN training for the binarization task.**

code detection. The binarized output will be passed onto a standard barcode reader, as reported in the next section. We evaluate how the combination of our ML model and barcode reader enables data storage capabilities in Section 6.

**5.3.2 Real-Time Inference.** We developed a *Unity3D* application that interfaces with both the phone’s standard RGB camera and our IR camera module in order to both extract the embedded data and display it in a visually appealing manner by means of AR augmentation of the tracked environment.

The application works by constantly running the ML model on the feed from the IR camera and then redirecting the output to a code reader library *Dynamsoft SDK*<sup>17</sup> for extraction. The output from the reader consists of a list of the values of all detected QR codes and their coordinates. In the case of offline embedding, we concatenate all results from the read QR codes in order to reconstruct the encoded string. In the case of online embedding, each QR code corresponds to a location on the data server that allows the application to download further instructions.

## 6 Technical Evaluation

In this section, we will evaluate various factors affecting the IR ink machine detectability and data capacity, such as IR ink density, detection distance and QR code’s error correction code (ECC). In the following, we explain the effects of each variable. Starting with increasing the ECC, designers can improve the machine detectability of the QR codes, but increasing the ECC too much requires additional paper space for the same embedding. On the other hand,

by increasing the distance between the NIR camera and the IR content, the illuminance also decreases, and thus so does the machine detectability of the content. Finally, by increasing the ink density, we increase the contrast, and thus also the machine detectability. However, increasing the ink density might cause it to become visible to the human eye, and since we already determined the lower limits of human perception in Section 4, the amount of IR ink that can be used in IMPRINTO documents is capped.


Therefore, we examine the effects of detection distance and ECC at invisible IR ink density levels on the data capacity of the sheet as it is crucial to better understand the trade-offs of using invisible IR content. As Data Matrix codes do not provide the option to set an ECC, we focus on QR codes in our evaluation. Our results can guide design decisions when fabricating IMPRINTO documents.

### 6.1 Data Capacity Based on IR Ink Density, Detection Distance and ECC

**6.1.1 Procedure.** For the purposes of this evaluation, we will focus on the three invisible ink densities associated with the background color classes outlined earlier in this paper: 81%, 102% and 155% (see Figure 9a). We focus on the distances necessary to capture a US letter-size sheet (25.5 cm in order to capture an 8.5" x 11" sheet) and a US half letter-size sheet (19.5 cm in order to capture a 5.5" x 8.5" sheet) as they present common real-world scenarios. Please note that our current setup features an NIR camera that has a slightly larger field of view than the area illuminated by the integrated LEDs. However, for the purpose of evaluating the maximum capacity, we will assume that the whole sheet is uniformly illuminated.

<sup>17</sup><https://www.dynamsoft.com/>





**Figure 10: Examples of the samples used for bit size evaluation at 25.5 cm camera distance necessary for letter-size sheets.**

After capturing the images, we use our image processing pipeline consisting of the ML model we trained (Section 5.3.1) and the code reader (Section 5.3.2) to evaluate the results.

**Smallest Detectable Module.** In order to determine the maximum data capacity, we must assess the size of the smallest detectable module [17]. Our sample printing precision was capped at 300 DPI, thus we were able to only assess the size of the smallest detectable module to within 0.085 mm (1" / 300 DPI). For each configuration of distance, ink density, and ECC (low, medium, and high), we printed a sample containing 12 randomly generated QR codes, that were not part of the training set used in Section 5.3.1, ranging in total module size from 0.67 mm to 1.6 mm in increments of 0.085 mm arranged in a 3 by 4 grid (see Figure 10). This range was determined based on empirical trials working with QR code detection after the initial training. In a real-world setting, each QR code would be exposed in multiple frames and from several slightly different positions. In order to better reflect this in our evaluation, we aligned the camera with each of the four columns in the grid and captured a picture for each column. We define the smallest detectable module as the smallest one that was detectable in at least one of the four pictures. Thus, our estimates for maximum data capacity only consider QR codes that can be detected. An example of the data for two such samples can be seen in Figure 10.

**Maximum Data Capacity.** Similarly to our perceptual experiment, our maximum data capacity estimates are conservative, since we enforce a quiet zone of at least four modules around each QR code in order to comply with the *ISO/IEC 8859-1* standard. In this first iteration of our work, we assume that we cover the whole sheet with QR codes of identical size instead of trying a more complex packing algorithm that may involve mixing QR codes of different sizes. We included QR codes with varying data capacities, version 1 to version 7<sup>18</sup>, because in version 7 we reached our largest module size with a quiet zone of 4 modules. Under the assumption that all QR codes are identical in size, we can iterate through all QR code versions (1–7) and try to fit them into a rectangular pattern to maximize the number of characters.

**6.1.2 Results.** The key results of our data capacity evaluation can be found in Figure 11 and Table A 2. Here, we report estimates for

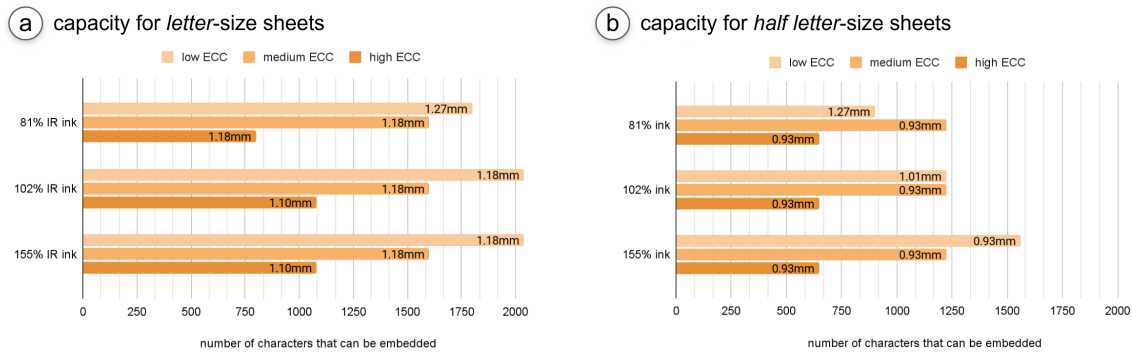
the binary encoding scheme of QR codes. In cases where the data is limited to numeric or alphanumeric characters, we can increase capacity by leveraging the QR code standard’s native support.

**Letter-size.** The maximum capacity results for our letter-size sheets show that for all three IR ink densities, we achieve similar maximum capacities, with the smallest detectable module size of 1.18 mm under the medium and low ECC. 81% IR ink seems to fall short of the other two classes with a smallest detectable module size of 1.27 mm. This is to be expected since its contrast is rather low, making it the most challenging case. ECC showed the expected effect: a reduction in data capacity but improved detection, which can be seen as a further validation of our approach.

**Half letter-size.** The findings for our half letter-size sheets suggest mixed results. First, the data capacities seem promising, given that only half of the space is available. The improved detection can most likely be attributed to the closer distance of the camera to the sheet. This is also reflected in the smallest detectable module size of 0.93 mm, which is significantly smaller than in the letter-size condition. We want to highlight that we did not train our model described in the previous section on this distance, but it can still detect QR codes reliably. However, despite 155%, the difference between low and medium ECC appears to have negligible effects.

**Comparison to traditional binarization approaches.** To assess the effectiveness of the resulting CNN-based approach, we also compared it to traditional binarization methods used in previous work [17, 19], which applies both contrast limited adaptive histogram equalization (CLAHE) [80] and adaptive thresholding [64], passing both binarized versions separately to the QR code reader. We used the same dataset from the previous section to determine how effectively the CNN binarization can outperform the traditional filter-based image processing approach. Our approach resulted in a 102.7% increase in the total number of QR codes detected compared to the previous method. Notably, the CNN makes a larger impact on the detection of low-IR ink (Class 1) samples compared to high-IR ink (Class 3), i.e., an increase of 164.0% vs. 92.7%. This is due to the fact that the pipeline for traditional methods vary based on factors such as lighting and the degree of code distortion, and may thus need constant adjustment of filter parameters. In contrast, our CNN-based approach is more scalable, given its training dataset was crafted for our specific goal.

<sup>18</sup><https://www.qrcode.com/>



**Figure 11: The effect of IR ink amount and error correction code on the data capacity and module size of the sheet.**

Together, our results provide users with estimates of how much content can be embedded using our IR ink watermarking technique.

## 7 Use Cases and Future Applications

In this section, we explore various applications of our watermarking technology designed for quick and integrated paper augmentation. We aim to demonstrate the breadth of use cases that IMPRINTO can support and additional benefits that our method provides to users (e.g., invisible, mobile, privacy-preserving) across various contexts. Here, we present applications that are unique to IMPRINTO, and we use replication for further validation [46].

### 7.1 Embedding "Paper DNA"

IMPRINTO could allow users to embed the digital metadata of the document as an intrinsic part of the physical paper. Utilizing this intrinsic "paper DNA," users could enable further authentication solutions.

**7.1.1 Recreating Documents from Fragments.** The "Paper DNA" watermark allows fragments of documents to be scanned and linked back to the full digital copy. For example, if a poster for an educational event is partially torn, as shown in Figure 12a, a student can scan the remaining portion to retrieve the original poster and access the missing information. The advantage of IMPRINTO is that reconstruction can work even if the document is heavily fragmented. This is possible due to the high data capacity of a single marker that is sufficient to recreate the document. Because the information is repeatedly encoded without influencing the document's aesthetics, only one small part is necessary to retrieve its entire content. Maia et al. demonstrated this for 3D printed objects, where redundancy renders the information readable from a broken or damaged object to recover its original model [51], which inspired our application. For paper materials, image-based tracking cannot easily achieve this because it would require training on fragmented documents, which does not scale well. We envision this feature can also benefit marketing strategies by including tearable sections in posters or ads with embedded links for future reference, which are known to occupy too small space to include extra information, and IMPRINTO can overcome these physical space constraints.

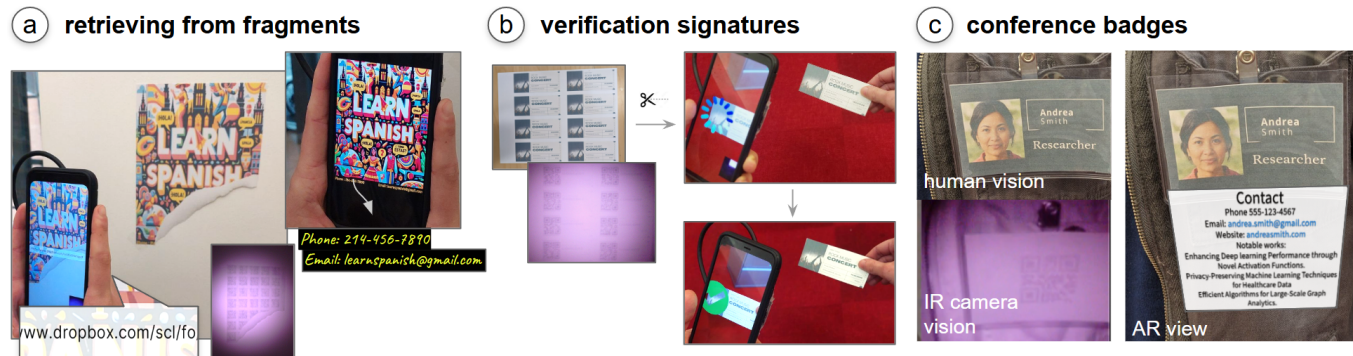
**7.1.2 Imprinted Verification Signatures.** IMPRINTO enhances document authenticity by embedding verification signatures in printed materials in combination with other techniques. For example, printed contracts could include invisible blockchain-based signatures or hyperlinks to the digital version for comparison against post-print modifications. If doubts arise, scanning with IMPRINTO can confirm the document's originality. The advantage is that this adds a second layer of "invisible" security by preventing people with malicious intent from reading names, passcodes, or signatures directly from the tickets, e.g., by shoulder surfing. This application exploits humans' unawareness of the presence of invisible IR content.

Figure 12b shows an example where a friend group bought multiple tickets to a concert. The individual tickets, after being cut out from the sheet, each contain a verification code, which is scanned at the entrance of the venue.

### 7.2 Augmenting Documents with Contextual Annotations and Multimedia

IMPRINTO seamlessly integrates digital interactivity into traditional paper media, from educational enhancements to professional networking solutions.

**7.2.1 Smart Business Cards and Conference Badges.** IMPRINTO could be used for fabricating intelligent business cards or conference badges as envisioned by Want et al. [74], which could, for example, also enable access to restricted areas, digital contact exchanges, and enhance networking experiences at professional gatherings without any electronic components. For instance, at conferences, attendees could wear badges and exchange business cards enhanced with IMPRINTO, as shown in Figure 12c. Scanning a badge reveals the attendee's professional profile, contact information, and links to their work, facilitating networking without the need for physical contact or exchanging multiple items. This works best if attendees know that all badges or business cards support IR content—thus, they know how to retrieve this information. Additionally, they can hide information they do not want to share immediately with anyone who passes by because it is directly visible on the badge.



**Figure 12: IMPRINTO is used for making digital metadata an intrinsic part of the printed document to (a) recreate documents from fragments or (b) verify the authenticity of individual pieces. (c) The application augments normal vision by adding the AR content, providing additional information that would overload the limited space of a conference badge.**

**7.2.2 Interactive Flashcards for Learning.** IMPRINTO enhances traditional flashcards by integrating multimedia, aiding tasks like language learning. We replicate certain functionalities of CV-based flashcards, which include pronunciation guides via simple scans, turning materials into interactive experiences that improve outcomes [52]. However, our application also supports offline access to hints and answers, minimizing online distractions. Students can reveal hints or check answers through an app, promoting independent, self-paced learning. Flashcards have very restricted space, severely limiting the kinds of tracking markers that can be embedded. IMPRINTO uses the entire space of the flashcard, i.e., front and back, even without any content, because it does not interfere with it. Nevertheless, image or text detection could be used to achieve the same. However, we would expect IMPRINTO to perform better during occlusive hand interactions with the flashcards since we can embed a larger number of evenly distributed markers.

**7.2.3 Offline High-Capacity Data Access for Sensitive Topics.** IMPRINTO can discreetly embed critical information into materials, useful for raising awareness about stigmatized topics or promoting sensitive products (e.g., healthcare) while ensuring privacy. For instance, a health organization can distribute posters on mental health with embedded AR markers, allowing access to additional information, such as treatment options, via a simple scan, as shown in Figure 1c, without leaving a digital footprint. This demonstrates IMPRINTO's key contribution: enabling invisible, high-capacity offline data embedding without requiring users to constantly send data to third-party servers. We also see potential in using offline high-capacity data access to enhance the accessibility of printed media, such as alt-texts, to images for visually impaired people without adapting visual designs to accommodate markers.

**IMPRINTO's Utility.** The examples in this section illustrate how IR watermarking can function as the "back end" for interactions with printed documents. Much like in web development or software engineering, this approach emphasizes a clear separation between the presentation layer (the "front end") and the data access layer (the "back end"). By abstracting and concealing technical, computer-readable components, our system offers a seamless, user-friendly interface while maintaining robust functionality. This separation

enables new possibilities for interaction, surpassing the limitations of visual markers or standard QR codes in terms of abstraction and user experience. We argue that, if IR watermarks become the "back end" of printed media one day, it may pave the way for transformative applications without compromising the complexity and depth of the underlying technology.

### 7.3 Origami: Interactive Folding Instructions

IMPRINTO brings a twist to the traditional art of origami by embedding interactive, step-by-step folding instructions directly onto the paper. As users progress with their folds, new instructions appear on the exposed surfaces, as shown in Figure 1b, based on the number and combination of the IMPRINTO codes visible to the IR camera. The advantage of IMPRINTO is that even if the origami paper has no visible content (i.e., white sheet) in order not to impact the look of the final artifact, the sheet can still carry invisible folding instructions.


### 7.4 Sustainable On-Demand Watermarking

IMPRINTO empowers users to watermark existing posters and documents larger than letter-size format on demand (see Figure 1a). This may help to save resources by avoiding re-printing posters or documents in case of missing content. Traditionally, people often go through multiple iterations until they arrive at their desired design. IMPRINTO enables people to focus on visual aesthetics with the possibility of adding content wherever and whenever needed.

### 7.5 Beyond Paper Augmentation

By using a mobile printer, IMPRINTO could also add watermarks to plastic [53] (a), fabric [22] (b), and wood [28] (c) and various form factors demonstrated in Figure 13. This is possible because the NIR ink is dye-based and can, therefore, be absorbed by these materials. In addition, our CNN detection pipeline can be fine-tuned for detecting, e.g., QR codes on convex or concave surfaces, highlighting the breadth of our contribution. This opens up new possibilities for embedding invisible identification markers to personal belongings (a) and (b). In Figure 13c, we also show how IMPRINTO may be used for robotics use cases, in this case for robotic vacuum cleaners to safely avoid the pet's area.





**Figure 13: Envisioned extension of IMPRINTO. (a) A user wears AR smart glasses that capture IMPRINTO content by exploiting the device's built-in always-on NIR cameras. (b, c, d) Watermarking on different materials and surfaces to identify personal belongings and facilitate robotic navigation through invisible cues.**

Through these scenarios, IMPRINTO shows its potential of integrating infrared inkjet watermarking into everyday interactions, and enabling solutions across educational, professional, and personal contexts.

## 8 Discussion

The value of hybrid paper-digital interfaces that augment physical documents with digital content has been well established [29]. To realize hybrid documents, we propose a novel watermarking method that uses IR-based printing to store digital content as an intrinsic part of the document. Unlike previous approaches that rely on a network connection and external storage for the digital assets, our approach with IMPRINTO maintains the document format as the single container of both physical and digital contents. In this section, we discuss the limitations of using IMPRINTO from the perspectives of authoring, printing, detection, and consumption. We also discuss future work to improve the approach.

### 8.1 Towards Multimodal Content

IMPRINTO's interface currently supports only a rudimentary form of AR that only permits 2D visual content and audio and does not leverage multimodal content to its full potential. We hope that we will enable authors to embed more intricate objects and interactions in the future. Adding support for multimedia assets will pave the way to expanding on previously examined use cases [1, 60]. The most impactful addition would be that of 3D assets since that would allow users to utilize the document as a spatial anchor.

### 8.2 Effects on Printed IR Ink

Our experiments investigated human and machine detectability under controlled indoor lighting and viewing conditions. However, the perception and reliability of printed IR ink can vary significantly depending on external factors such as illumination levels, viewing angles, and the type of paper used. One key limitation is that low-light conditions may reduce the effectiveness of human detectability due to insufficient IR reflection, while excessive illumination, such as direct sunlight, could lead to overexposure, affecting both machine readability and ink longevity. Future work could explore adaptive imaging techniques or optimized illumination setups to mitigate these issues. Regarding *paper type*, we anticipate that glossy or coated paper may introduce reflection

artifacts, which could interfere with machine-based recognition systems by creating specular highlights that obscure ink contrast. Conversely, highly porous or rough-textured paper might cause ink diffusion, potentially reducing print sharpness and affecting detection accuracy. A promising direction is to systematically evaluate these effects using an experimental framework similar to Xu et al.'s methodology, where QR code robustness was tested under varying lighting conditions, scanning angles, and environmental factors [78]. Applying a similar protocol could help quantify how different conditions impact IR ink detection and visibility.

As with any other dye-based ink, UV-dependent fading is an inherent limitation [57] that is important to consider for document reliability and permanence. Constant exposure to heavy UV sources such as the sun can cause the inks to fade over time. Willis et al. [76] showed how different IR ink types can be used in conjunction with UV-resistant coatings to achieve 98% contrast preservation under office lighting conditions. We envision that for commercial use cases where long-term preservation is desired, a similar coating can be applied, which is sufficient for indoor applications.

### 8.3 Invisible IR Ink

We used a psychophysical experiment to determine conservative estimates for the IR ink densities that remain invisible to users for a wide range of different background colors. As with most experiments, we were limited by the number of background colors that we could include in the experiment. Therefore, we do not know how our results generalize to other background colors, combinations of colors, and different graphic patterns. Nevertheless, we developed a system that facilitates embedding with invisible IR ink for any RGB color. Note that our method is very conservative and favors invisibility over machine detection. The actual DTs are likely much higher, which should further improve machine detectability. Nevertheless, we contribute a methodology that designers can apply to determine the "sweet spot" between visibility, data capacity, and machine detectability.

We also want to highlight that the gradient IR ink and background color bars used in our psychophysical experiment differ from QR codes which may have had an effect on our estimated invisible IR ink DTs. We decided against using QR codes in the experiment because (1) it would have only been possible to test QR codes at predefined densities (e.g., at 20%, 40%, ...), limiting the precision of the study, and (2) the limited number of samples, which

would not reflect the variability of QR codes of different sizes. As a result, our experiment also aims to provide general estimates for shapes other than QR codes and may, therefore, be used as a starting point for any type of invisible IR-printed marker.

## 8.4 Discoverability and Practicality

To demonstrate IMPRINTO, we used NIR-based fabrication and detection tools. While NIR cameras are getting popular in many handheld devices (e.g., *iPhone* and *iPad* use it for facial recognition and LIDAR 3D scanning), not all platforms currently give 3rd-party developers access to the raw NIR stream (e.g., only the processed depth map can be accessed on *iOS*). We argue the interest in NIR applications will increase as more use cases are demonstrated by future projects.

In our current implementation, IMPRINTO documents can optionally be marked with a small icon or visual label on the document in our embedding tool to allow users to discover embedded AR content. We envision that next-generation AR hardware can more fully leverage IMPRINTO's utility. Compared to using a handheld device, always-on AR smart glasses such as *Meta Orion* could be constantly scanning the environment for hidden AR content [5, 9]. Currently, most AR glasses already leverage NIR cameras, but mainly for localizing the user and mapping their environment for 3D tracking purposes. With head-worn AR glasses with integrated NIR cameras used for IMPRINTO sensing, the user would not have to manually capture objects using the phone during the interaction. We recommend future research to focus on implementing this on AR glasses.

## 9 Conclusion

We introduced IMPRINTO, a watermarking technique that embeds computer-readable information while remaining invisible to the human eye. Unlike previous methods, our approach utilizes the entire document, including white spaces, and works regardless of background color. Through a psychophysical experiment, we determined the maximum ink that can be embedded without being detected. We developed tools to support users in applying IR ink technology, including software for efficient information embedding and a universal camera module for capturing IMPRINTO watermarks. Our open-source ML pipeline processes these images for robust use with standard QR code readers. We demonstrated various use cases, highlighting the potential of invisible IR content for hybrid paper-digital interfaces and advancing watermarking techniques.

## References

- [1] Andrea Alessandrini, Alessandro Cappelletti, and Massimo Zancanaro. 2014. Audio-augmented paper for therapy and educational intervention for children with autistic spectrum disorder. *International Journal of Human-Computer Studies* 72, 4 (2014), 422–430. Publisher: Elsevier.
- [2] Ananta Narayanan Balaji, Clayton Kimber, David Li, Shengzhi Wu, Ruofei Du, and David Kim. 2023. RetroSphere: Self-Contained Passive 3D Controller Tracking for Augmented Reality. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 4 (Jan. 2023), 157:1–157:36. <https://doi.org/10.1145/3569479>
- [3] Steve Benford, Boriana Koleva, Anthony Quinn, Emily-Clare Thorn, Kevin Glover, William Preston, Adrian Hazzard, Stefan Rennick-Egglestone, Chris Greenhalgh, and Richard Mortier. 2017. Crafting Interactive Decoration. *ACM Trans. Comput.-Hum. Interact.* 24, 4, Article 26 (aug 2017), 39 pages. <https://doi.org/10.1145/3058552>
- [4] Leah Buechley, Sue Hendrix, and Mike Eisenberg. 2009. Paints, paper, and programs: first steps toward the computational sketchbook. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*. ACM, Cambridge United Kingdom, 9–12. <https://doi.org/10.1145/1517664.1517670>
- [5] Daniel Campos Zamora, Mustafa Doga Dogan, Alexa F Siu, Eunye Koh, and Chang Xiao. 2024. MoiréWidgets: High-Precision, Passive Tangible Interfaces via Moiré Effect. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/3613904.3642734>
- [6] Zhutian Chen, Wai Tong, Qianwen Wang, Benjamin Bach, and Huamin Qu. 2020. Augmenting Static Visualizations with PapARVis Designer. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–12. <https://doi.org/10.1145/3313831.3376436>
- [7] Tingyu Cheng, Koya Narumi, Youngwook Do, Yang Zhang, Tung D. Ta, Takuya Sasatani, Eric Markvicka, Yoshihiro Kawahara, Lining Yao, Gregory D. Abowd, and HyunJoo Oh. 2020. Silver Tape: Inkjet-Printed Circuits Peeled-and-Transferred on Versatile Substrates. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 4, 1 (March 2020), 1–17. <https://doi.org/10.1145/3381013>
- [8] Enrico Costanza and Jeffrey Huang. 2009. Designable Visual Markers (*CHI '09*). Association for Computing Machinery, New York, NY, USA, 1879–1888. <https://doi.org/10.1145/1518701.1518990>
- [9] Francesco Riccardo Di Gioia, Eugenie Brasier, Emmanuel Pietriga, and Caroline Appert. 2022. Investigating the Use of AR Glasses for Content Annotation on Mobile Devices. *Proceedings of the ACM on Human-Computer Interaction* 6, ISS (Nov. 2022), 574:430–574:447. <https://doi.org/10.1145/3567727>
- [10] Mustafa Doga Dogan. 2024. Ubiquitous Metadata: Design and Fabrication of Embedded Markers for Real-World Object Identification and Interaction. <https://doi.org/10.48550/arXiv.2407.11748> arXiv:2407.11748 [cs].
- [11] Mustafa Doga Dogan, Patrick Baudisch, Hrvoje Benko, Michael Nebeling, Huaishu Peng, Valkyrie Savage, and Stefanie Mueller. 2022. Fabricate It or Render It? Digital Fabrication vs. Virtual Reality for Creating Objects Instantly. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 5. <https://doi.org/10.1145/3491101.3516510>
- [12] Mustafa Doga Dogan, Vivian Hsinyueh Chan, Richard Qi, Grace Tang, Thijs Roumen, and Stefanie Mueller. 2023. StructCode: Leveraging Fabrication Artifacts to Store Data in Laser-Cut Objects. In *Proceedings of the 8th ACM Symposium on Computational Fabrication (SCF '23)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3623263.3623353>
- [13] Mustafa Doga Dogan, Steven Vidal Acevedo Colon, Varnika Sinha, Kaan Aksit, and Stefanie Mueller. 2021. SensiCut: Material-Aware Laser Cutting Using Speckle Sensing and Deep Learning. In *Proceedings of the 34th Annual ACM Symposium on User Interface Software and Technology*. ACM, Virtual Event USA, 15. <https://doi.org/10.1145/3472749.3474733>
- [14] Mustafa Doga Dogan, Faraz Faruqi, Andrew Day Churchill, Kenneth Friedman, Leon Cheng, Sriram Subramanian, and Stefanie Mueller. 2020. G-ID: Identifying 3D Prints Using Slicing Parameters. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3313831.3376202>
- [15] Mustafa Doga Dogan, Raul Garcia-Martin, Patrick William Haertel, Jamison John O'Keefe, Ahmad Taka, Akarsh Aurora, Raul Sanchez-Reillo, and Stefanie Mueller. 2023. BrightMarker: 3D Printed Fluorescent Markers for Object Tracking. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST '23)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3586183.3606758>
- [16] Mustafa Doga Dogan, Eric J Gonzalez, Karan Ahuja, Ruofei Du, Andrea Colaco, Johnny Lee, Mar Gonzalez-Franco, and David Kim. 2024. Augmented Object Intelligence with XR-Objects. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. Association for Computing Machinery, New York, NY, USA, 1–15. <https://doi.org/10.1145/3654777.3676379>
- [17] Mustafa Doga Dogan, Alexa F. Siu, Jennifer Healey, Curtis Wigington, Chang Xiao, and Tong Sun. 2023. StandARone: Infrared-Watermarked Documents as Portable Containers of AR Interaction and Personalization. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, 1–7. <https://doi.org/10.1145/3544549.3585905>
- [18] Mustafa Doga Dogan, Ahmad Taka, Michael Lu, Yunyi Zhu, Akshat Kumar, Aakar Gupta, and Stefanie Mueller. 2022. InfraredTags: Embedding Invisible AR Markers and Barcodes Using Low-Cost, Infrared-Based 3D Printing and Imaging Tools. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New Orleans LA USA, 9. <https://doi.org/10.1145/3491102.3501951>
- [19] Mustafa Doga Dogan, Veerapatr Yotamornsunthorn, Ahmad Taka, Yunyi Zhu, Aakar Gupta, and Stefanie Mueller. 2022. Demonstrating InfraredTags: Decoding Invisible 3D Printed Tags with Convolutional Neural Networks. In *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, 7. <https://doi.org/10.1145/3491101.3519905>
- [20] Fatih Dogangun, Serdar Bahar, Yigit Yildirim, Bora Toprak Temir, Emre Ugur, and Mustafa Doga Dogan. 2024. RAMP: Robotic Augmented Reality for Machine Programming and Automation. <https://doi.org/10.48550/arXiv.2410.13412> arXiv:2410.13412 [cs].



- [21] Tobias Drey, Jessica Janek, Josef Lang, Dietmar Puschmann, Michael Rietzler, and Enrico Rukzio. 2022. SpARklingPaper: Enhancing Common Pen- And Paper-Based Handwriting Training for Children by Digitally Augmenting Papers Using a Tablet Screen. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 6, 3 (Sept. 2022), 113:1–113:29. <https://doi.org/10.1145/3550337>
- [22] Jack Forman, Mustafa Doga Dogan, Hamilton Forsythe, and Hiroshi Ishii. 2020. DefeXtiles: 3D Printing Quasi-Woven Fabric via Under-Extrusion. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. ACM, Virtual Event USA, 1222–1233. <https://doi.org/10.1145/3379337.3415876>
- [23] Jiayun Fu, Bin Zhu, Weiwei Cui, Song Ge, Yun Wang, Haidong Zhang, He Huang, Yuanyuan Tang, Dongmei Zhang, and Xiaojing Ma. 2021. Chartem: Reviving Chart Images with Data Embedding. *IEEE Transactions on Visualization and Computer Graphics* 27, 02 (Feb. 2021), 337–346. <https://doi.org/10.1109/TVCG.2020.3030351> Publisher: IEEE Computer Society.
- [24] Sergio Garrido-Jurado, Rafael Muñoz-Salinas, Francisco José Madrid-Cuevas, and Manuel Jesús Marín-Jiménez. 2014. Automatic Generation and Detection of Highly Reliable Fiducial Markers under Occlusion. *Pattern Recognition* 47, 6 (2014), 2280–2292.
- [25] Christopher Getschmann and Florian Echtler. 2021. Seedmarkers: Embeddable Markers for Physical Objects. In *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction (TEI '21)*. Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3430524.3440645>
- [26] Lisa Gitelman. 2014. *Paper knowledge: Toward a media history of documents*. Duke University Press.
- [27] Nan-Wei Gong, Jürgen Steimle, Simon Olberding, Steve Hodges, Nicholas Edward Gillian, Yoshihiro Kawahara, and Joseph A. Paradiso. 2014. PrintSense: a versatile sensing technique to support multimodal flexible surface interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, Toronto Ontario Canada, 1407–1410. <https://doi.org/10.1145/2556288.2557173>
- [28] Xinyue Gui, Ding Xia, Wang Gao, Mustafa Doga Dogan, Maria Larsson, and Takeo Igarashi. 2025. Draw2Cut: Direct On-Material Annotations for CNC Milling. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA.
- [29] Feng Han, Yifei Cheng, Megan Strachan, and Xiaojuan Ma. 2021. Hybrid Paper-Digital Interfaces: A Systematic Literature Review. In *Designing Interactive Systems Conference 2021 (DIS '21)*. Association for Computing Machinery, New York, NY, USA, 1087–1100. <https://doi.org/10.1145/3461778.3462059>
- [30] Kaimeing He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep Residual Learning for Image Recognition. <https://doi.org/10.48550/arXiv.1512.03385> arXiv:1512.03385 [cs].
- [31] David Holman, Roel Vertegaal, Mark Altosaar, Nikolaus Troje, and Derek Johns. 2005. Paper windows: interaction techniques for digital paper. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 591–599.
- [32] Jeremy Howard and Sylvain Gugger. 2020. fastai: A Layered API for Deep Learning. *Information* 11, 2 (Feb. 2020), 108. <https://doi.org/10.3390/info11020108> arXiv:2002.04688 [cs, stat].
- [33] Ramya Iyer, Mustafa Doga Dogan, Maria Larsson, and Takeo Igarashi. 2025. XR-penter: Material-Aware and In Situ Design of Scrap Wood Assemblies. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA.
- [34] Sam Jacoby and Leah Buechley. 2013. Drawing the electric: storytelling with conductive ink. In *Proceedings of the 12th International Conference on Interaction Design and Children (IDC '13)*. Association for Computing Machinery, New York, NY, USA, 265–268. <https://doi.org/10.1145/2485760.2485790>
- [35] Joshua D.A. Jung, Rahul N. Iyer, and Daniel Vogel. 2019. Automating the Intentional Encoding of Human-Designable Markers. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland Uk) (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3290605.3300417>
- [36] Joshua D.A. Jung, Rahul N. Iyer, and Daniel Vogel. 2019. Automating the Intentional Encoding of Human-Designable Markers. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland Uk) (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–11. <https://doi.org/10.1145/3290605.3300417>
- [37] Y. Kakehi, T. Hosomi, M. Iida, T. Naemura, and M. Matsushita. 2006. Transparent Tabletop Interface for Multiple Users on Lumisight Table. In *First IEEE International Workshop on Horizontal Interactive Human-Computer Systems (TABLETOP '06)*. 8 pp.–. <https://doi.org/10.1109/TABLETOP.2006.34>
- [38] Henry R. Kang. 1994. Gray component replacement using color mixing models. In *Color Hard Copy and Graphic Arts III*, Vol. 2171. SPIE, 287–296. <https://doi.org/10.1117/12.175317>
- [39] Hirokazu Kato and Mark Billinghurst. 1999. Marker Tracking and HMD Calibration for a Video-based Augmented Reality Conferencing System. In *Proceedings 2nd IEEE and ACM International Workshop on Augmented Reality (IWAR'99)*. IEEE, 85–94.
- [40] Yoshihiro Kawahara, Steve Hodges, Benjamin S. Cook, Cheng Zhang, and Gregory D. Abowd. 2013. Instant inkjet circuits: lab-based inkjet printing to support rapid prototyping of UbiComp devices. In *Proceedings of the 2013 ACM international joint conference on Pervasive and ubiquitous computing*. ACM, Zurich Switzerland, 363–372. <https://doi.org/10.1145/2493432.2493486>
- [41] Han-Jong Kim, Ju-Whan Kim, and Tek-Jin Nam. 2016. miniStudio: Designers' Tool for Prototyping UbiComp Space with Interactive Miniature. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 213–224. <https://doi.org/10.1145/2858036.2858180>
- [42] Frederick A. A. Kingdom and Nicolaas Prins. 2016. Chapter 5 - Adaptive Methods. In *Psychophysics (Second Edition)*. Academic Press, San Diego, 119–148. <https://doi.org/10.1016/B978-0-12-407156-8.00005-0>
- [43] Diederik P. Kingma and Jimmy Ba. 2017. Adam: A Method for Stochastic Optimization. <https://doi.org/10.48550/arXiv.1412.6980> arXiv:1412.6980 [cs].
- [44] Konstantin Klamka, Wolfgang Büschel, and Raimund Dachselt. 2017. Illuminated Interactive Paper with Multiple Input Modalities for Form Filling Applications. In *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces (ISS '17)*. Association for Computing Machinery, New York, NY, USA, 434–437. <https://doi.org/10.1145/3132272.3132287>
- [45] Konstantin Klamka and Raimund Dachselt. 2017. IllumiPaper: Illuminated Interactive Paper. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 5605–5618. <https://doi.org/10.1145/3025453.3025525>
- [46] David Ledo, Steven Houben, Jo Vermeulen, Nicolai Marquardt, Lora Oehlberg, and Saul Greenberg. 2018. Evaluation Strategies for HCI Toolkit Research. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (Montreal QC, Canada) (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–17. <https://doi.org/10.1145/3173574.3173610>
- [47] Johnny Lee, Scott Hudson, and Pau Dietz. 2007. Hybrid infrared and visible light projection for location tracking. In *Proceedings of the 20th annual ACM symposium on User interface software and technology (UIST '07)*. Association for Computing Machinery, New York, NY, USA, 57–60. <https://doi.org/10.1145/1294211.1294222>
- [48] Marjorie R. Leek. 2001. Adaptive Procedures in Psychophysical Research. *Perception & Psychophysics* 63, 8 (Nov. 2001), 1279–1292. <https://doi.org/10.3758/BF03194543>
- [49] Chao Li, Cai Yin Wang, and Shu Jie Wang. 2013. A Black Generation Method for Black Ink Hiding Infrared Security Image. *Applied Mechanics and Materials* 262 (2013), 9–12. Publisher: Trans Tech Publ.
- [50] Zhen Li, Michelle Annett, Ken Hincley, Karan Singh, and Daniel Wigdor. 2019. HoloDoc: Enabling Mixed Reality Workspaces that Harness Physical and Digital Content. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, Glasgow Scotland Uk, 1–14. <https://doi.org/10.1145/3290605.3300917>
- [51] Henrique Teles Maia, Dingzeyu Li, Yuan Yang, and Changxi Zheng. 2019. Layer-Code: optical barcodes for 3D printed shapes. *ACM Transactions on Graphics* 38, 4 (July 2019), 112:1–112:14. <https://doi.org/10.1145/3306346.3322960>
- [52] Kher Hui Ng, Hai Huang, and Claire O'malley. 2018. Treasure Codes: Augmenting Learning from Physical Museum Exhibits through Treasure Hunting. *Personal and Ubiquitous Computing* 22 (2018), 739–750.
- [53] Mehmet Ozdemir, Marwa AlAlawi, Mustafa Doga Dogan, Jose Francisco Martinez Castro, Stefanie Mueller, and Zjenja Doubrovski. 2024. Speed-Modulated Ironing: High-Resolution Shade and Texture Gradients in Single-Material 3D Printing. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3654777.3676456>
- [54] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc. [https://proceedings.neurips.cc/paper\\_files/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html](https://proceedings.neurips.cc/paper_files/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html)
- [55] Narjes Pourjafarian, Marion Koelle, Fjolla Mjaku, Paul Strohmeier, and Jürgen Steimle. 2022. Print-A-Sketch: A Handheld Printer for Physical Sketching of Circuits and Sensors on Everyday Surfaces. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (New Orleans, LA, USA) (CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 270, 17 pages. <https://doi.org/10.1145/3491102.3502074>
- [56] William Preston, Steve Benford, Emily-Clare Thorn, Boriana Koleva, Stefan Rennick-Egglestone, Richard Mortier, Anthony Quinn, John Stell, and Michael Worboys. 2017. Enabling Hand-Crafted Visual Markers at Scale. In *Proceedings of the 2017 Conference on Designing Interactive Systems (Edinburgh, United Kingdom) (DIS '17)*. Association for Computing Machinery, New York, NY, USA, 1227–1237. <https://doi.org/10.1145/3064663.3064746>
- [57] MaxMax Llewellyn Data Processing. 2022. IR Ink. <https://maxmax.com/phosphordyesandinks/infrared-phosphors-dyes-and-inks/infrared-down-conversion-powder/ir-ink-down-conversion>



- [58] Jie Qi and Leah Buechley. 2010. Electronic popables: exploring paper-based computing through an interactive pop-up book. In *Proceedings of the fourth international conference on Tangible, embedded, and embodied interaction*. ACM, Cambridge Massachusetts USA, 121–128. <https://doi.org/10.1145/1709886.1709909>
- [59] Jing Qian, Qi Sun, Curtis Wigington, Han L. Han, Tong Sun, Jennifer Healey, James Tompkin, and Jeff Huang. 2022. Dually Noted: Layout-Aware Annotations with Smartphone Augmented Reality. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–15. <https://doi.org/10.1145/3491102.3502026>
- [60] Shwetha Rajaram and Michael Nebeling. 2022. Paper Trail: An Immersive Authoring System for Augmented Reality Instructional Experiences. In *CHI Conference on Human Factors in Computing Systems*. ACM, New Orleans LA USA, 1–16. <https://doi.org/10.1145/3491102.3517486>
- [61] Chris Ridpath and Wendy Chisholm. 2000. *Techniques For Accessibility Evaluation And Repair Tools*. Technical Report. W3C. <https://www.w3.org/TR/AERT/#color-contrast>
- [62] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. <https://doi.org/10.48550/arXiv.1505.04597> arXiv:1505.04597 [cs].
- [63] Daniela K. Rosner and Kimiko Ryokai. 2010. Spyn: augmenting the creative and communicative potential of craft. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 2407–2416. <https://doi.org/10.1145/1753326.1753691>
- [64] J. Sauvola and M. Pietikäinen. 2000. Adaptive document image binarization. *Pattern Recognition* 33, 2 (Feb. 2000), 225–236. [https://doi.org/10.1016/S0031-3203\(99\)00055-2](https://doi.org/10.1016/S0031-3203(99)00055-2)
- [65] Chen Song, Zhengxiong Li, Wenyao Xu, Chi Zhou, Zhanpeng Jin, and Kui Ren. 2018. My Smartphone Recognizes Genuine QR Codes!: Practical Unclonable QR Code via 3D Printing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2 (July 2018), 1–20. <https://doi.org/10.1145/3214286>
- [66] Hyunyoung Song, Tovi Grossman, George Fitzmaurice, François Guimbretiere, Azam Khan, Ramtin Attar, and Gordon Kurtenbach. 2009. PenLight: combining a mobile projector and a digital pen for dynamic visual overlay. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 143–152.
- [67] Martin Spindler, Sophie Stellmach, and Raimund Dachselt. 2009. PaperLens: advanced magic lens interaction above the tabletop. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*. ACM, Banff Alberta Canada, 69–76. <https://doi.org/10.1145/1731903.1731920>
- [68] Jürgen Steimle, Mohammadreza Khalilbeigi, Max Mühlhäuser, and James D. Hollan. 2010. Physical and digital media usage patterns on interactive tabletop surfaces. In *ACM International Conference on Interactive Tabletops and Surfaces*. ACM, Saarbrücken Germany, 167–176. <https://doi.org/10.1145/1936652.1936685>
- [69] Matthew Tancik, Ben Mildenhall, and Ren Ng. 2020. StegaStamp: Invisible Hyperlinks in Physical Photographs. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Seattle, WA, USA, 2114–2123. <https://doi.org/10.1109/CVPR42600.2020.00219>
- [70] Daniel Wagner, Gerhard Reitmayr, Alessandro Mulloni, Tom Drummond, and Dieter Schmalstieg. 2008. Pose Tracking from Natural Features on Mobile Phones. In *2008 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*. 125–134. <https://doi.org/10.1109/ISMAR.2008.4637338>
- [71] Daniel Wagner, Gerhard Reitmayr, Alessandro Mulloni, Tom Drummond, and Dieter Schmalstieg. 2010. Real-Time Detection and Tracking for Augmented Reality on Mobile Phones. *IEEE Transactions on Visualization and Computer Graphics* 16, 3 (2010), 355–368. <https://doi.org/10.1109/TVCG.2009.99>
- [72] Hsi-Chun Wang, Wen-Hsin Liu, Chia-Long Chang, and Yi-Hui Chen. 2008. Design of Halftone-Based AR Markers under Infrared Detection. In *2008 International Conference on Computer Science and Software Engineering*, Vol. 6. 97–100. <https://doi.org/10.1109/CSSE.2008.1391>
- [73] Zeyu Wang, Shiyu Qiu, Qingyang Chen, Natallia Trayan, Alexander Ringlein, Julie Dorsey, and Holly Rushmeier. 2019. AniCode: authoring coded artifacts for network-free personalized animations. *The Visual Computer* 35, 6–8 (June 2019), 885–897. <https://doi.org/10.1007/s00371-019-01681-y>
- [74] Roy Want, Kenneth P. Fishkin, Anuj Gujar, and Beverly L. Harrison. 1999. Bridging physical and virtual worlds with electronic tags. In *Proceedings of the SIGCHI conference on Human factors in computing systems the CHI is the limit - CHI '99*. ACM Press, Pittsburgh, Pennsylvania, United States, 370–377. <https://doi.org/10.1145/302979.303111>
- [75] Denso Wave. 2022. Information capacity and versions of QR Code. <https://www.qrcode.com/en/about/version.html>
- [76] Karl D. D. Willis, Takaaki Shiratori, and Moshe Mahler. 2013. HideOut: mobile projector interaction with tangible objects and surfaces. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction (TEI '13)*. Association for Computing Machinery, New York, NY, USA, 331–338. <https://doi.org/10.1145/2460625.2460682>
- [77] Chang Xiao, Cheng Zhang, and Changxi Zheng. 2018. FontCode: Embedding Information in Text Documents Using Glyph Perturbation. *ACM Transactions on Graphics* 37, 2 (April 2018), 1–16. <https://doi.org/10.1145/3152823>
- [78] Mingliang Xu, Qingfeng Li, Jianwei Niu, Hao Su, Xiting Liu, Weiwei Xu, Pei Lv, Bing Zhou, and Yi Yang. 2021. ART-UP: A Novel Method for Generating Scanning-Robust Aesthetic QR Codes. *ACM Trans. Multimedia Comput. Commun. Appl.* 17, 1 (April 2021), 25:1–25:23. <https://doi.org/10.1145/3418214>
- [79] Peiying Zhang, Chenhui Li, and Changbo Wang. 2021. VisCode: Embedding Information in Visualization Images using Encoder-Decoder Network. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (Feb. 2021), 326–336. <https://doi.org/10.1109/TVCG.2020.3030343> Conference Name: IEEE Transactions on Visualization and Computer Graphics.
- [80] Karel Zuiderveld. 1994. Contrast limited adaptive histogram equalization. In *Graphics gems IV*. Academic Press Professional, Inc., USA, 474–485.

## A Appendix

We provide a table for background color selection for the experiment and another table regarding the smallest module size and character capacity comparison.

Group ID	Darkness (analog $K$ value)	Calculated luminescence ( $lum$ ) range	Number of selected colors (roughly four colors per 25% $lum$ .)
1	[0%, 25%]	(8.55%, 100%]	16
2	[25%, 50%]	(5.7%, 75%]	12
3	[50%, 75%]	(2.85%, 50%]	8
4	[75%, 100%]	(0%, 25%]	4
5	100%	0%	1
6	Four colors have been added to enhance color hue coverage.		

**Table 1: Color selection table for the experiment. Colors for each tested group are chosen based on their darkness and luminescence ranges.**

Ink Type	Letter size (distance $\geq 25.5\text{cm}$ )		Half letter size (distance $\geq 19.5\text{cm}$ )	
	Smallest size	Maximum capacity	Smallest size	Maximum capacity
Low EEC 81% ink	1.27mm	1800 characters	1.27mm	900 characters
Medium EEC 81% ink	1.18mm	1600 characters	0.93mm	1224 characters
High EEC 81% ink	1.18mm	800 characters	0.93mm	648 characters
Low EEC 102% ink	1.18mm	2040 characters	1.01mm	1224 characters
Medium EEC 102% ink	1.18mm	1600 characters	0.93mm	1224 characters
High EEC 102% ink	1.10mm	1080 characters	0.93mm	648 characters
Low EEC 155% ink	1.18mm	2040 characters	0.93mm	1560 characters
Medium EEC 155% ink	1.18mm	1600 characters	0.93mm	1224 characters
High EEC 155% ink	1.10mm	1080 characters	0.93mm	648 characters

**Table 2: Comparison of module sizes and capacities at different ink levels and distances.**