Citation for published version:

Benedetti, L, Winnemöller, H, Corsini, M & Scopigno, R 2014, Painting with Bob: Assisted creativity for novices. in *UIST 2014 - Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, pp. 419-428, 27th Annual ACM Symposium on User Interface Software and Technology, UIST 2014, Honolulu, UK United Kingdom, 5/10/14. <https://doi.org/10.1145/2642918.2647415>

DOI:

[10.1145/2642918.2647415](https://doi.org/10.1145/2642918.2647415)

Publication date:

2014

Document Version

Early version, also known as pre-print

[Link to publication](#)

© ACM, 2014. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in *IST 2014 - Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*. ACM, pp. 419-428, 27th Annual ACM Symposium on User Interface Software and Technology, UIST 2014, Honolulu, United Kingdom, 5-8 October. <http://doi.acm.org/10.1145/2642918.2647415>

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Painting with Bob: Assisted Creativity for Novices

Luca Benedetti ISTI-CNR and University of Bath Bath, U.K. l.benedetti@bath.ac.uk	Holger Winnemöller Adobe Systems, Inc. Seattle, U.S.A. hwinnemo@adobe.com	Massimiliano Corsini ISTI-CNR Pisa, Italy m.corsini@isti.cnr.it	Roberto Scopigno ISTI-CNR Pisa, Italy r.scopigno@isti.cnr.it
---	---	---	--

Figure 1. Several examples of geese painted by novice users with *Painting with Bob* in under ten minutes.

ABSTRACT

Current digital painting tools are primarily targeted at professionals and are often overwhelmingly complex for use by novices. At the same time, simpler tools may not invoke the user creatively, or are limited to plain styles that lack visual sophistication. There are many people who are not art professionals, yet would like to partake in digital creative expression. Challenges and rewards for novices differ greatly from those for professionals. In this paper, we leverage existing works in Creativity and Creativity Support Tools (CST) to formulate design goals specifically for digital art creation tools for novices. We implemented these goals within a digital painting system, called *Painting with Bob*. We evaluate the efficacy of the design and our prototype with a user study, and we find that users are highly satisfied with the user experience, as well as the paintings created with our system.

Author Keywords

Painting; novices; creativity;

ACM Classification Keywords

I.3.4 Graphics Utilities: Paint systems; J.5 ARTS AND HUMANITIES: Fine arts; I.3.3 Picture/Image Generation

INTRODUCTION

“Every child is an artist, the problem is staying an artist when you grow up” – Pablo Picasso

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UIST 2014, October 5–8, 2014, Honolulu, HI, USA.

Copyright © 2014 ACM 978-1-4503-3069-5/14/10 ...\$15.00.

<http://dx.doi.org/10.1145/2642918.2647415>

Most children readily engage in many forms of creative expression, including play, pretending, and art projects such as painting and drawing [14, 15]. But as we grow older, many adults disengage from creative expression, with explanations such as “*I cannot draw*”, “*I am not creative*”, or “*I do not know what to draw*”. Various factors affect creativity and may explain this different approach between children and adults. Amabile [1] differentiates between self-motivated (pleasure) versus extrinsic (reward) motivations for the respective groups, and states that extrinsic motivations are often associated with anxieties that inhibit creativity. Time constraints in our busy daily lives may also adversely affect creativity [2]. Yet, despite all these obstacles, the benefits of (adult) creative expression are well documented [12].

Painting is a good example of this rift between childhood and adult creativity. Finger painting is accessible even to a young toddler, while the “proper” mastery of brushes and paints takes years of training. In this paper, we design a creativity support tool (CST) for digital painting for the vast majority of adults who rarely or never engage in fine arts. Artistic challenges and rewards are different for novices and professionals [9, 26]. The average novice does not seek the approval of an art critic or a paying customer, but rather seeks the enjoyment of the task and the approval of his (novice) peers. Novices face various challenges: They are unfamiliar with tools and techniques, they are intimidated by the blank page, they are unwilling to invest a lot of time, and they are easily frustrated or discouraged [9, 30]. We therefore seek to create a *safe* playground for novices to experiment with painting, where safe means “controlled to guarantee a minimum degree of quality”. Controlling the user’s environment necessarily limits their creative choices, so a balance must be struck between the two. The type of creativity we enable in this paper is that of a ‘cooking challenge’, where the ingredients and preparation techniques are pre-determined to maxi-

mize the likelihood of a pleasant-tasting dish, while the user is free to experiment within the constraints of the challenge.

Our painting system, affectionately dubbed “*Painting with Bob*” (PwB) is specifically designed to overcome the above artistic, technical, and creative challenges. We start by formulating a set of system design guidelines for creative support tools for novices. We then describe the implementation of PwB to fulfill the design requirements. We propose a practical user study to demonstrate the efficacy of our system design and report on its results. We see the main contributions of our paper as the design guidelines for a fine art CST for novices, along with a detailed case study of creating and evaluating a CST based on these guidelines. While we are clear that the results we obtained are specific to our digital painting tool (PwB), we argue that many of the design decisions and evaluation methods make no limiting reference to digital painting, and could be applied to other digital fine arts for novices, such as drawing, sketching, cartooning, etc.

RELATED WORK

Creativity is a complex phenomenon with varying definitions, depending on the context in which it is discussed. Using Beghetto and Kauffman’s [21] taxonomy, we focus on little-c and mini-c (personal level) creativity, rather than Big-C or Pro-c (domain level) creativity. Davis et al. [9] discuss the challenges of CSTs for novices. They distinguish two types of novices: The *tool* novice already has skills applicable to a given domain (e.g. a traditional painter), but an appropriate CST may increase their productivity (e.g. acquiring digital painting skills). In contrast, a *domain* novice is unfamiliar with both the domain as well as the CSTs used in that domain. Davis et al. illustrate this concept with a *skill-threshold* that needs to be overcome by the domain novice to even participate in that domain. The authors argue that since novices are often unwilling to invest sufficient time in learning, they are underrepresented within the CST community, yet require significantly different strategies to support their creativity [9]. Designing CSTs that artificially (through the help of computer assistance) lift the novice above the skill threshold can enable them to immediately participate in creative exploration, and that is the approach taken in this paper.

Much of research into Creativity Support Tools [30] focuses on professionals, as these most commonly engage in creative endeavors [26]. Creative professionals are exposed to significant external pressures (time, money, novelty, recognition), whereas novices are generally motivated by aspects of play [15, 28] and what Csikszentmihalyi calls “*flow*” [7], defined as “a mental state of complete immersion and energized focus”. While flow is not a prerequisite for creativity, people often experience flow during phases of extreme productivity or creativity. Specifically, flow is associated with the following subjective experiences: 1) There are clear goals every step of the way; 2) There is immediate feedback to one’s actions; 3) There is a balance between challenges and skills; 4) Action and awareness are merged; 5) Distractions are excluded from consciousness; 6) There is no worry of failure; 7) Self-consciousness disappears; 8) The sense of time becomes distorted; 9) The activity becomes autotelic (the purpose of the

activity is within itself). Within these concepts, our goals for a CST for novice art creation focus primarily on designing a user experience that supports all elements of flow, thereby maximizing the likelihood of such an experience for the user. That is, we define success in terms of a subjectively rewarding experience for the novice user, rather than some more objective measures of artistic mastery, novelty, or critical praise that might be applied to professionals. Our approach to CST design and evaluation is therefore different from those supporting professional or domain expert creativity [24].

Of those CSTs designed for novices, many focus on educational goals [29, 32]. iCanDraw [10] guides users through a series of steps to draw faces, providing automatic user feedback about the accuracy of their efforts. Iarussi et al. [19] presented an educational tool to practice traditional drawing-by-observation techniques. Cummings et al. [8] developed a tutorial system to draw eyes. While those works demonstrated skill improvement in their user studies, the actual tutorials are highly specific and do not scale to different artistic skills or motifs. Sketch-Sketch Revolution by Fernquist et al. [11] addressed the scalability issue by basing their tutorials on recorded sessions of real artists. In contrast to these works, our goal is *not* to teach users traditional skills, as this is necessarily time consuming and delays gratification, thereby creating an impedance for engagement. Instead, we aim to enable novices to partake in the artistic process immediately, even if this means that initially they can only be successful within the constraints of the CST. If an individual feels encouraged by their experience to acquire the necessary skills to substitute for the CST, we consider that an added bonus, but we do not see it as a requirement of our approach.

Non-educational CSTs can be classified as supporting *tool* or *domain* novices [9]. Kazi et al. [22]’s *Vignette* system helps with the design and manipulation of textures in pen-and-ink illustration. Their approach simplifies an artist’s workflow but still requires significant skills and must thus be considered a *tool novice* CST. A notable example of a *domain novice* CST, much aligned with our own motivations, is the *ShadowDraw* system by Lee et al. [23]. Here, a database of images is analyzed for strong contours, matched to user input, and blended to create “shadows” of many possible forms, allowing for serendipitous exploration of the objects in the database. While the authors demonstrate interesting artworks by novices, the paper focuses on the technical aspects of the system and omits user-centric aspects such as subjective satisfaction, ownership, personal style, etc., all of which we believe to be important design goals for domain novice CSTs.

Several commercial systems bear superficial similarities to PwB. *PsykoPaint* is a web application that allows the user to upload an image and paint over it. Painting tools are grouped in “styles” with many free parameters for each style and the associated usability and discoverability issues. *GMX PhotoPainter* uses image analysis to create smart tools (similar to PwB), however their analysis seems limited, requiring a large number of free parameters. *Studio Artist* replaces many free parameters with “thousands of presets”, which can only be discovered through exhaustive experimentation. We argue

that thousands of presets are just as unmanageable for novices as dozens of parameters. Microsoft's *Freshpaint*, implements recent oilpaint techniques [5] and provides a paint-by-number mode for novices, but apart from visual tracing lines the user still has to paint mostly unassisted. *Artrage* is a media simulation system only offering automatic color selection, similar to the "simple" assistance mode we present in this paper. *Photoshop* is a well-known CST. Its standard brushing model contains no assistive mode and many dozens of parameters. Photoshop's *Art History Brush* tool allows a user to paint stylized strokes based on image content. This tool itself is less sophisticated than the above systems in terms of visual quality, automated assistance, and ease of use. For these reasons, and to give a better comparison between varying degrees of assistance in CSTs for novices, we compare PwB with a simplified version of itself, rather than with existing commercial systems.

For rendering, our paper takes inspiration from earlier efforts on digital painting based on stroke libraries and orientation fields [27]. However, in those works the user interaction is limited to parameter selection (e.g. [17]) or semantic disambiguation of the image components (e.g. [35]).

SYSTEM DESIGN

In this section, we review several creativity frameworks, and we distill design goals from these theories that pertain to our specific problem domain of art creation CSTs for novices.

Schneiderman [30] suggests four design principles for general CSTs: S1) support exploratory research; S2) enable collaboration; S3) provide rich history-keeping; and S4) design with "low thresholds, high ceilings, and wide walls". Of these, we find only S4 to be relevant in our context. Specifically, *Low thresholds* suggests a minimal training requirement, while *High ceilings* represent support for levels of sophistication (e.g. from first-timer to casual enthusiast). Finally, *Wide walls* describes the breadth of expression of a CST (e.g. personal style).

Read et al. [25] propose three dimensions of *fun*. R1) The perceived experience is better than the predicted one; R2) The user feels engrossed in the experience; and R3) A willingness exists to continue or repeat an activity.

Rubin et al. [28] suggest six factors of *play*, many of which overlap the nine elements of *flow* [7], listed again in shortened form: F1) Clear goals; F2) Immediate feedback; F3) Balanced challenges; F4) Merged action and awareness; F5) No distractions; F6) No fear; F7) No self-consciousness; F8) Ignorance of time; F9) Autolectic activity.

In selecting our design goals for art CSTs for novices we draw from several of the aforementioned theories (listed in parentheses), but we also add novel goals that are specific to our intended domain and audience:

N1) Kickstart Novices often do not know what to create ("blank page problem"). The CST should offer a simple and obvious starting point.

N2) Easy to learn Effectively needing little or no training, thereby reducing the time investment necessary to engage in the creative process (S4,F3).

N3) Easy to use The UI should be simple and intuitive. The user should easily find the right tool, and the tool's response to user actions should be predictable (F4).

N4) WYSIWYG The user should be given immediate and accurate visual feedback about the results of their interactions with the system (F2).

N5) Guaranteed success Aesthetically pleasing results (to the user) should be attainable in a short amount of time, with possible improvements for greater time investment (S4,F6).

N6) Ownership & Achievement Despite any technological assistance, users should identify with having created the final artifact. Ideally, they should feel proud of their creation (F7).

N7) Enjoyment The user should find enjoyment in the act of working with the system (R1,R2,R3,F9)

N8) Creative flexibility User guidance and free-form flexibility are opposing goals. Within the confines of the assistive technology, the user should feel as much freedom for creative expression as possible (F3).

N9) Individual results Different users should be able to create different results. Allowing users to develop a personal style is a vital aspect of creative expression.

The above design goals are useful for a variety of art creation CSTs for novices. However, the act of *painting* that we have chosen for our case study, carries some domain-specific challenges in support of the general goals, above:

P1) Brushes Common digital brushing systems have an overwhelming number of free parameters that are difficult for the novice to adjust → PwB brushes should be simple to configure, yet be visually appealing with a range of natural media appearances, such as watercolor or acrylic (N2,N3,N5,N8).

P2) Technique A typical painting consists of hundreds or thousands of individual brush strokes. The domain novice does not have the knowledge or skill to apply individual strokes for the desired effect → PwB should offer the user control at a higher level than individual strokes, thus reducing the task from painting each stroke to painting more or less "detail" (N2,N3,N8,N9).

P3) Fidelity Skilled artists perform brushes strokes that, in aggregate, suggest texture, materials, etc. → PwB must use image analysis techniques to automatically choose and configure brushes to reproduce accurate local detail (N3,N5).

Other domain-specific techniques exist, such as *Dark-to-light*, *Coarse-to-fine*, and *Back-to-front* (see Results section). A question arises as to which of these techniques should be addressed by the CST. Our approach is minimalistic: We want to empower novices without restricting them too much by overly constraining the user experience (N8). Referring to Davis et al.'s [9] "*skill-threshold*", we choose to include the techniques that we deem necessary to lift a user above the threshold, whereas we purposely exclude techniques which

merely increase productivity beyond that level. That is, we offer the user enough support to engage with a creative domain (e.g. P1-P3), but we leave it to the user to acquire any additional high-level skills (Dark-to-light, etc.)

IMPLEMENTATION

Fundamentally, PwB is a simple system, related to Haeblerli's [16] original work. PwB offers the user a virtual canvas and a set of tools to paint on it. The user starts by loading a source image. The system analyzes the image and computes various metrics. The tools may query any metric to augment the user interactions, thereby producing "smart" tool behavior. The user interface (Fig. 2, Left) and experience are kept minimalistic (N2, N3, N4). The user can choose between four tools and use these to paint over a faded version of the source image. Using an existing image to paint on addresses goal N1. The only user-specified tool parameters are tool-size and brush-alpha. Additionally, the user can pan/zoom the painting, adjust the visibility of the reference image and undo/redo actions. The system is a multi-threaded C++ application, uses Lua scripting for tool behavior, OpenCV for image analysis, and OpenGL for rendering.

Brushes – Paint marks are chosen from a brush library of 56 post-processed scans of real brush strokes. First, we painted several dozen paint strokes (long/round, straight/wavy, dab/smear, dry/wet) with green acrylic paint on white paper. We scanned the marks and separated them in Photoshop into two channels: a texture channel representing bristle details, and an alpha channel representing the shape and opacity of paint (Fig. 2, Right). A paint mark can be rendered as-is, or be deformed along a Beziér path. A global brush alpha value can be increased to create a watercolor-like appearance, or decreased to simulate thicker paint. A fragment shader, based on the brush's texture channel, adds shadows and specular highlights. PwB's brushing system is fast, flexible, and produces realistic looking brush strokes, thereby addressing (by design) goals N2, N3, N5, P1, P2, and P3.

Analysis – To drive smart tool behavior (N3,N5,P3), we compute the following analysis measures: *Meanshift* filtering [6] simplifies the source image's color palette for sampling by tools. Successive Gaussian filter passes are applied for sampling by larger tool radii. *Edge Tangent Flow (ETF)* [20] computes a set of progressively smoothed orientation fields. The *Laplacian response* of the grayscale image, indicates regions of high frequency content. *CIE Y variance* approximates local saliency. *Oriented Gabor filter* responses [13] indicate local isotropy and orientation histograms. Image analysis is performed on downsampled ($\sim 1\text{Mpix}$) source images at load-time, takes ~ 20 seconds (unoptimized code), and is cached for future sessions.

User Actions – As the user brushes over the canvas, the cursor coordinates are regularized by fitting them to a spline which is then sampled according to tool size. Each resampled position holds coordinates and input data, such as screen-space speed, and pressure, tilt & rotation, where available.

A tool's Lua script uses this data, along with any data queried from the analysis metrics, to issue paint commands to the ren-

dering system. Each command consists of: texture-id, position, orientation, size, color, transparency, texture highlights and shadows, and optional deformation along a spline. A tool script can be viewed as a mapping of user input and image metrics, onto rendering commands. Interactivity is maintained by multi-threaded computation of tool scripts and rendering, thereby addressing N4, and facilitating N6-N9.

Tool Design – To avoid confusing the user with too many tools, we empirically determined a minimal set of tools. We first created a large number of brush behaviors based on physical brushes and painting tutorials. We then painted many images in the subject categories that we wanted PwB to support: portraits, rural landscapes, cityscapes, and nature. We compared notes on which brushes were useful in creating the visual aspects of each motif. The foliage and noisy textures of rural landscapes and nature scenes suggested a brush that quickly fills large areas with many paint marks. The regular structure of hair, fur, and architecture suggested a brush with long strokes along such structures, etc. Based on our observations, we distilled a set of four tools: (1) The *Single* tool paints a sequence of individual paint marks; (2) The *Fill* tool covers an area with a number of marks; (3) The *Structure* tool paints long, possibly curved paint strokes; and (4) the *Eraser* tool removes paint.

Simple & Smart Tools – To test the effect that painting assistance would have on the quality of images produced with PwB (N5) and user experience (N6-N9) we created two versions of the four tools, a *simple* version with little assistance, and a *smart* version with more assistance (Fig. 3).

The *simple toolset* automatically chooses a (size-dependent) color and contrast values (highlights & shadows) for each brush mark. Additionally: (1) The *Simple Single* tool paints a sequence of paint marks following the input path. The user's speed controls the brush's aspect ratio and spacing along the path; (2) The *Simple Fill* fills the area under the cursor with a number of simple single instances; (3) The *Simple Structure* creates a single long stroke along the user's path with the color of the path's starting point; and (4) The *Simple Eraser* deletes all paint within the cursor area.

The *smart toolset* enables higher-level user control by adding assistance beyond the simple toolset: (1) The *Smart Single* tool aligns paint marks along the (size-dependent) ETF orientation field and uses local Gabor responses to adjust brush rendering: In low contrast, isotropic regions, brush textures near unit aspect-ratio are chosen and rendered large and more transparent, with low brush contrast. Paint marks in high-contrast and anisotropic regions are chosen to be small, oblong, more opaque, with high brush contrast; (2) The *Smart Fill* applies many smart single instances within the cursor area. The tool implements a local coarse-to-fine technique by sorting selected paint marks by aspect ratio, rendering them in order of increasing value; (3) The *Smart Structure* tool generates a series of splined brushes, each following local ETF orientations. Splines are sub-divided at significant color discontinuities. Brush transparency and contrast are modulated as above; (4) The *Smart Eraser* deletes only areas that have a color similar to the starting point of the user path.

Figure 2. (Left) The user interface, showing tools, commands, UI elements (zoom and sliders), views, and dual-cursor. (Right) Brush components.

Figure 3. Examples of a single user interaction. Top: Simple tools. Bottom: Smart tools. From left-to-right: Single, Fill, Structure, and Eraser

EVALUATION

To demonstrate the efficacy of our PwB implementation in supporting the general goals for art creation for novices (N1-N9) and specifically for digital painting (P1-P3) we propose a targeted user study. Specifically, we are interested in two main questions: (1) *Is the design of PwB successful in achieving the aforementioned goals?* (2) *Is there a substantial impact of user assistance through smart brushes on these goals?*

Verification and analysis of CSTs is a complex topic, partly due to the complex nature of creativity itself. Sternberg [31] argues that different types of creativity may be used in acts of drawing versus writing versus problem solving, and that different strategies are needed to measure each type. Carroll et al. [3] list three basic approaches to measuring creativity used by the community, along with caveats and benefits: (1) *Self-reporting* is a cheap and easy analysis tool, commonly taking the form of a survey. However, problems with self-evaluation for task-performance and other objective quantities are widely acknowledged within the community [18]. (2) *Psychophysical Measures* include biometrics like EEG, fMRI, galvanic skin response, etc. They provide the most objective data, but are expensive and difficult to procure, calibrate, and operate. (3) *External Judges* may be able to offer a reasonably objective opinion about the quality of a given work. Challenges include the cost and time-requirements of such domain experts. Hocevar [18] also cautions that judges may not be able to distinguish between technical skills, aesthetics, and creativity.

Due to the shortcomings of each method, Carroll et al. [3] perform all three methods, and triangulate the results to im-

prove overall confidence of the findings. While the authors observed interesting correlations between pairs of the three methods, they acknowledged that the setup and analysis costs are prohibitive and only increased confidence slightly.

Our own study design is driven by three main goals: (1) *Relevance*: As discussed above we are not interested in traditional aspects of creativity for professionals. Rather, we want to promote creative experiences for novices. Hocevar [18] asserts that most practical experiments measure creativity *correlates*, rather than creativity itself. The correlates we are interested in are related to *flow*; (2) *Reproducibility*: We see our work as a first step in a more general framework for CSTs for novice art creation. Thus, we want our methodology to apply to similar CSTs, and be affordable to implement and evaluate; (3) *Reliability*: Despite the above constraints, we want the results of our user study to make quantifiable statements about the success and failure of a specific CST.

Like Carroll et al. [3] we propose a multi-pronged approach to measure the multitude of facets affecting creativity. We instrumented PwB to record all UI interactions, allowing us to create usage statistics for the application, and representing a de facto screen recording of each user session, which could then be coded and analyzed. Unlike a screen recording, our usage logs can be queried and analyzed electronically. As several of our design goals are highly subjective (e.g. N6,N7), we employed self-reporting to obtain relevant data. Finally, we used external judges to rate paintings created with PwB. Since our ultimate goal is not high art, and the most likely audience of novice paintings are personal social circles, we recruited peer judges instead of domain experts.

Assistance Levels – In our user study, we compare tools with two levels of painting assistance (*simple* vs. *smart*) to measure the effect of assistance on our design goals. We specifically do not compare with commercial tools to limit biases due to UI, overall software complexity, and brush quality.

Protocol – We selected 20 volunteers from over 50 candidates, to balance gender, favor minimal digital painting experience, non-computer scientists, and a broad age range. Participants were offered a voucher for their efforts.

Each user took part in two sessions of about one hour each. We randomly assigned users to two groups to control for mode sequence: Group A used *simple* tools in the first session and *smart* tools in the second session. Group B did the reverse. To minimize painting strategies from the first session influencing the second session, we spaced the sessions 1-4 days apart (depending on subject availability). Each session was subdivided in the following phases: Introduction→Learning→Painting→Feedback.

The *Introduction* phase of the *first* session began by explaining the aim of the study and its phases. Then an intake questionnaire was administered (Q1).

The *learning* phase was sub-divided into three steps: learning of the UI and tools (by showing a 5 minute video), and practical training (5–7 minutes). Volunteers were encouraged to ask questions during this step, as no feedback would be given during the painting phase (to prevent experimenter bias).

In the *painting* phase, users were asked to perform 3 paintings, with a maximum time of 10 minutes for each one. To allow for a variety of painting results while minimizing total session durations, the 3 input images were chosen by rotation from the 4 images shown in Fig. 5, which in turn were selected to represent the 4 supported subject categories (landscape, portrait, city, nature). Participants were encouraged to share their thoughts with the experimenter. Users could stop the painting once they were satisfied with the result.

The *final* phase consisted of a feedback questionnaire (Q2).

For the *introduction* phase of the *second* session, subjects were given the same introductions as before, without the questionnaire. The *learning* and *painting* phases were administered as before with the same 3 source images, in the same order, but with a new tool set.

In the *feedback* phase, subjects filled out questionnaire Q3. Finally, participants ranked the 6 paintings produced in both session with a visual sorting program. The initial sequencing of images was randomized and no meta-data from the images was shown.

Questionnaires – For brevity, the questionnaires are summarized here. Full versions are available online.

Intake questionnaire (Q1) gathers demographic (gender, age, normal/corrected vision), as well as experience and skill assessments (computer, traditional art, digital art).

Questionnaire (Q2), presented after each user session, asks questions related to design goals (N1-N9) and domain chal-

lenges (P1-P3). The questions are partially derived from the Creativity Support Index (CSI) [4], but adapted to a domain novice CST. For example questions, see Table 1.

Questionnaire Q3, consists of two parts. The first part is identical to Q2 (for the latest session). Part two asks about personal preferences between the two tool sets used, the amount of control over the tools, the perceived creativity, and the subjective quality of the produced paintings.

Logs – All user sessions were recorded with detailed event logs and 5-second-interval screenshots. The logs captured all paint events, UI adjustments (alpha, tool size, zoom, etc.) as well as command usage (undo, redo, clear), so that exact statistics could be computed about usage patterns.

One researcher (not physically involved with the user study) coded usage behavior, as follows: Scanning all session recordings in random order, common vs. distinct usage patterns (features) were identified. In a second pass, again randomized, feature occurrences were noted for each session.

Peer Evaluation – We argue that peer judges (social circles) are the most likely critics of novice artwork and that they might use different evaluation criteria compared to expert judges. We thus asked 30 volunteers (not overlapping with those of the main study) to evaluate the quality of paintings produced with PwB. This allowed us to compute a relative ranking of images, users, and sessions.

RESULTS

Data of the three questionnaires was aggregated following the ITU-T Rec. P. 910 [33] by modeling scores with gain and bias (per subject) plus noise to identify systematic and random errors. An ANOVA analysis indicated that a gain and bias normalization was not required. A Kurtosis analysis (Annex 2 of [34]) was used to calculate range values to identify outliers, resulting in the removal of 2 of 20 subjects.

Data was aggregated by calculating the average (s_j) and standard deviations (σ_j) for each question, reported here as $s_j \pm \sigma_j$. For multiple choices questions, the data was aggregated by computing a histogram of the respective answers.

Population – Questionnaire Q1 tells us basic information about the demographic that participated in the user study. Subjects are divided almost equally between males (44%) and females (56%), and ages are distributed with 28% in the 15–30 range, 33% in the 30–40 range, and 39% in the 40 and above range. Most subjects had never taken art lessons (78%), and only 56% had decorated or painted something. While many subjects had prior experience with creative software (78% photo-editing, 61% drawing, 44% video-editing), they rated their skills in using such software as mediocre (photo editing: 5.0 ± 1.42 , drawing software: 4.0 ± 2.26 , all out of 10). As such, our average study subject was not entirely naïve in digital art creation, but could certainly be considered novice (non-professional), as per our intent.

Aggregate Data – Table 1 lists the results for feedback questions of Q2 and part one of Q3, regarding the subjective goals N5–N9. Average scores are high for many of the factors un-

Feedback (FB1-FB6)	Scores (10)
1-I like to play with this software (N7)	8.69 $\pm$ 1.15
2-I like the paintings I created (N5, N6, N7)	6.94 $\pm$ 1.52
3-I would share the paintings on my social network (N6, N9)	6.04 $\pm$ 2.74
4-This software makes me feel creative (N6, N9)	6.83 $\pm$ 2.17
5-This software is able to increase my creativity (N6, N9)	7.06 $\pm$ 1.84
6-I cannot realize these images without this software (N5)	8.64 $\pm$ 1.62

Table 1. Selected questions with associated goals and survey scores

der examination such as enjoyment, guaranteed success, creative flexibility etc. Two very high values include 8.7 ± 1.15 for FB1, and 8.6 ± 1.62 for FB6. The only score lower than we hoped for, with 6.0 ± 2.74 , is for FB3, aimed at judging ownership & achievement. However, while answering the questionnaire, most users spontaneously commented that this is because they could not choose the source image, and that they were pleased with their paintings. Feedback was quite uniform with two exceptions. Users felt more creative in their second session (7.2 ± 1.9) compared with the first (6.5 ± 2.3). This may be a sign that users were successfully lifted above the *skill-threshold* of creative engagement [9] during the first session. Users also liked paintings created with smart tools significantly more (7.2 ± 1.4) than those created with simple tools (6.6 ± 1.5). Additionally, 94.4% of users stated that they would like to have this software on their PC/smartphone, while 72.2% of subjects would pay 5 euros for a print of one of their paintings. Overall, the survey responses indicate a high degree of enjoyment and personal satisfaction with the user experience and artworks created. Note that similar results could be obtained with other CSTs. In fact, it is our hope that our paper might inspire comparative studies between multiple novice CSTs.

Simple vs. Smart tools – An important aspect of our user study was the degree to which user assistance (*simple* or *smart*) had an effect on the goals of our system design, particularly the interplay between assistance and creative freedom and perceived control.

As stated above, numerical ratings per image indicated that users preferred *smart tool* paintings. We obtain similar results when looking at user preference by session: 50% preferred paintings produced in smart tool sessions, vs. 39% who preferred paintings from simple tool sessions (11% stated no preference). In contrast, most people preferred the simple tool sessions compared to the smart sessions (50% vs. 33%), they experienced more control over the system (50% vs. 33%), and felt more creative (50% vs. 39%). As expected, there exists a strong correlation between control, and perceived flexibility and creativity. So while increased assistance leads to a higher level of satisfaction with the final paintings, this needs to be weighed against the need of users to feel “in charge”. Given the overall high feedback scores, it seems that PwB managed to strike a reasonable balance.

Log Analysis

While we instructed users in the use of the UI, we did not advise them on any painting related strategies. Consequently, when coding the users’ log files we observed a number of common painting technique challenges, along with tech-

Figure 4. (Left): Foreground elements (geese) drawn first; (Right): Attempts to fill in background (space between geese) results in irregular background and noticeable seams between foreground and background.

niques that users naturally (without guidance) developed to overcome these challenges.

Painter’s algorithm – The painter’s algorithm refers to a technique that draws background (BG) elements first, followed by foreground (FG) elements. This approach is beneficial as it naturally deals with the occlusions of BG elements by FG elements. This strategy seems unnatural to novices, as they tend to first paint what interests them most (generally FG elements). The users then have trouble painting a coherent BG without affecting FG elements. This often leads to notable seams between FG & BG areas (Fig. 4). While only one user demonstrated a proper back-to-front approach in the first image, about half the users developed this skill by the last painting of the second session.

Coarse-to-fine – This technique refers to directing the viewer’s attention by the amount of detail that is painted in different parts of the image. That is, a viewer looks where there is more amount of detail. This effect is commonly achieved by painting most of the image with coarse settings, and then refining areas that are semantically important. Our coding finds this technique to be discovered (not evident in early paintings, but present in later ones) by at least 6 users.

Content aware vs. Content agnostic vs. Novel content – Particularly for simple tools there were different approaches to brush movement over the image (N6,N8,N9). These approaches can be categorized as: (*Content aware*) User brushes emulating image content. E.g. swirly for clouds or wavy for water; (*Content agnostic*) User brushes unrelated to content, often in back-and-forth motions to quickly fill an area. Smart tools tend to produce better results for such behavior; (*Novel content*) User brushes in patterns unrelated to image content, but with an intent to create elements not part of the original image (add path to barn, smiley in the sky, etc.), a sign of playfulness and creativity (N7,N8,N9). We observed half the users paint content aware and half content agnostic in the beginning. At least 4 content agnostic users developed content aware skills by their last image. Notably, at least 7 users showed signs of novel content painting.

Preconceived notions

Just as most users lacked knowledge of proper painting techniques, some users’ strategies were clearly influenced by prior experiences in different media or other assumptions.

Sketcher’s technique – At least 7 users started by drawing thin outlines of objects and then filled these outlines with a

larger tool. We assume that novices more commonly sketch than paint, and thus try to transfer this skill. About half of the sketchers developed strategies more suitable to painting, while others were reasonably successful with their technique.

Paint on white assumption – We observed many users erasing a region before re-painting, whereas in most observed cases it would have been simpler to just layer new paint over old paint, without erasing. This reluctance to paint over existing paint might explain the challenges in discovering back-to-front or coarse-to-fine painting, which are best achieved by painting over existing paint.

Peer Judgments

Each of the 30 peer judges ranked the 4 best and worst paintings for each of the 4 source images. Despite not being expert judges, rankings were quite consistent: Votes for best images (by user) were 3.3 standard deviations from the mean, while votes for the worst paintings were 2.8 deviations from the mean. Of the best ranked images over all judges, 227 paintings were created with smart tools vs. 253 painted with simple tools, a slight but insignificant difference. However, for the worst paintings 175 were created with smart tools vs. 305 with simple tools. We interpret this finding to indicate that our smart tools do not help people create the best paintings at the high end, but they seem to create a minimum level of quality at the low end, compared with manual tools. This satisfies our goal N5 (*Guaranteed success*).

Additionally, we observe a definite learning effect: While 196 preferred paintings come from the first user session, 284 paintings come from the second one. Even within each session, the last paintings are rated higher on average than the first paintings, with most improvement noticed during the first session. While some of this learning effect could be increasing familiarity with handling of a stylus, it should be noted that PwB requires very little precision (e.g. compared to sketching). We also found evidence of skill improvements through the log analysis and therefore conclude that goal N2 (*Easy to learn*) is satisfied within PwB.

In retrospect, using the 4 best and worst pictures instead of a complete image ranking may introduce a potential bias due to outliers. A frequency analysis of our data shows that this bias is not present in our results, but we underline that a more reliable analysis should be based on a full rank.

DISCUSSION & FUTURE WORK

In this paper we proposed 9 general design goals for novice art CSTs (N1–N9), and 3 goals domain-specific to digital painting (P1–P3). The latter (P1–P3), N1 (Kickstart), and N4 (WYWISYG) were addressed by the technical implementation of PwB. The basic usage of PwB was *easy to use* (N3) and *easy to learn* (N2), as almost all users were able to produce pleasing looking paintings in under 10 minutes (low threshold). The average user improved their drawing skills dramatically within the first user session, with some slighter improvements in the second session, as rated by objective third party judges. More surprisingly, a number of users

developed basic painting techniques (coarse-to-fine, back-to-front, level-of-detail, etc.) without instructions, by mere experimentation (high ceiling). The goal of *guaranteed success* (N5) was achieved by multiple standards. The example images in this paper (e.g. Fig. 7) and the accompanying video demonstrate the visual quality of paintings produced with PwB. Of course, not every user was successful in obtaining high-quality images, as shown in Fig. 8. Satisfaction scores of users for their paintings were high, and most users stated that they would be unable to produce similar results without PwB. Responses to our *ownership and achievement* (N6) question in Q3 were lower than hoped, but partially explained by the lack of choice for the source images. *Enjoyment* (N7) ratings from the questionnaires were very high. One of our fears was that users would choose a default setting, fill the image carelessly and thus all generating a similar look. This was not the case. Each user was able to achieve *individual results* (N9), as demonstrated in Figs. 1 and 6. Finally, the assistance of the painting program should not unduly impede *creative flexibility* (N8). We were astonished by the unexpected and innovative paintings and techniques demonstrated by our users (e.g. Fig. 9), indicating that creative expression is possible even within a controlled system like PwB. An even higher level of creativity (compositional) could be achieved by allowing users to paint elements from different source images into a single painting.

One surprising result is that while *smart tools* increased subjective painting quality, the *simple tools* were still much appreciated for their sense of control. We believe that merely providing high-quality, yet simple-to-use brushes (auto-color; only 2 parameters, identical for every tool) could be sufficient to lift domain novices over the *skill-threshold*. In such case, future work might investigate more closely how this threshold is defined in terms of techniques or skill sets. Another surprise was the degree to which individuals naturally acquired various painting skills, with only minimal exposure to PwB (2 hours). We are interested in investigating how incidental learning could be maximized in novice art CSTs without negatively impacting our design goals. Finally, we would like to apply our design goals and evaluation techniques to other fine arts, such as sketching, cartooning, etc.

REFERENCES

1. Amabile, T. M. Motivation and creativity: Effects of motivational orientation on creative writers. *J. of Personality and Social Psych.* 48, 2 (1985), 393–399.
2. Amabile, T. M., Hadley, C. N., and Kramer, S. J. Creativity under the gun. *Harvard business review* 80, 8 (2002), 52.
3. Carroll, E. A., and Latulipe, C. Triangulating the Personal Creative Experience: Self-report, External Judgments, and Physiology. In *Proc. of Graphics Interfaces 2012* (2012), 53–60.
4. Carroll, E. A., Latulipe, C., Fung, R., and Terry, M. Creativity factor evaluation: Towards a standardized survey metric for creativity support. In *Proc. Creativity and Cognition '09*, ACM (2009), 127–136.

Figure 5. The four input images¹ for the user study. Each user painted three of these images.

Figure 6. An example of the diversity of final results obtained by different users from the same reference image.

Figure 7. Four “successful” results, where the system clearly allowed the user to express a well-defined style.

Figure 8. Four “failure” results, where the user was not able to exploit the system to reach the intended completion of the painting.

Figure 9. Four “creative” interesting results, where the user has been able to express a unique personal style.

5. Chu, N., Baxter, W., Wei, L.-Y., and Govindaraju, N. Detail-preserving paint modeling for 3D brushes. In *NPAR '10*, ACM (2010), 27–34.
6. Comaniciu, D., and Meer, P. Mean shift: a robust approach toward feature space analysis. *IEEE PAMI* 24, 5 (2002), 603–619.
7. Csikszentmihalyi, M. *Creativity: Invention, Flow and the psychology of discovery and invention*. Perennial & Harper, 1997.
8. Cummmings, D., Vides, F., and Hammond, T. I Don't Believe My Eyes!: Geometric Sketch Recognition for a Computer Art Tutorial. In *Proc. of SBIM'12* (2012), 97–106.
9. Davis, N., Winnemöller, H., Dontcheva, M., and Do, E. Y.-I. Toward a cognitive theory of creativity support. In *Proc. Creativity & Cognition '13* (2013), 13–22.
10. Dixon, D., Prasad, M., and Hammond, T. iCanDraw: Using Sketch Recognition and Corrective Feedback to Assist a User in Drawing Human Faces. In *Proc. SIGCHI '10*, ACM (2010), 897–906.
11. Fernquist, J., Grossman, T., and Fitzmaurice, G. Sketch-sketch revolution: an engaging tutorial system for guided sketching and application learning. In *UIST '11*, ACM (2011), 373–382.
12. Flood, M., and Phillips, K. D. Creativity in older adults: A plethora of possibilities. *Issues in Mental Health Nursing* 28, 4 (2007), 389–411.
13. Gabor, D. Theory of communication. part 1: The analysis of information. *Electrical Engineers-Part III: Radio and Communication Engineering, Journal of the Institution of* 93, 26 (1946), 429–441.
14. Gardner, H. *Art, Mind, And Brain: A Cognitive Approach To Creativity*. Basic Books, New York, NY, USA, 1982.
15. Goncu, A., and Gaskins, S. *Play and Development: Evolutionary, Sociocultural, and Functional Perspectives*. The Jean Piaget Symposium series. Taylor & Francis, New York, NY, USA, 2007.
16. Haeberli, P. Paint by numbers: Abstract image representations. *SIGGRAPH Comput. Graph.* 24, 4 (Sept. 1990), 207–214.
17. Hertzmann, A. Painterly rendering with curved brush strokes of multiple sizes. In *SIGGRAPH '98*, ACM (New York, NY, USA, 1998), 453–460.
18. Hocevar, D. Measurement of creativity: Review and critique. *Journal of Personality assessment* 45, 5 (1981), 450–464.
19. Iarussi, E., Bousseau, A., and Tsandilas, T. The drawing assistant: Automated drawing guidance and feedback from photographs. In *Proc. UIST '13*, ACM (New York, NY, USA, 2013), 183–192.
20. Kang, H., Lee, S., and Chui, C. K. Coherent line drawing. In *Proc. NPAR '07*, ACM (2007), 43–50.
21. Kaufman, J. C., and Beghetto, R. A. Beyond big and little: The four c model of creativity., 2009.
22. Kazi, R. H., Igarashi, T., Zhao, S., and Davis, R. Vignette: Interactive texture design and manipulation with freeform gestures for pen-and-ink illustration. In *Proc. CHI '12*, ACM (2012), 1727–1736.
23. Lee, Y. J., Zitnick, C. L., and Cohen, M. F. Shadowdraw: real-time user guidance for freehand drawing. *ACM Trans. Graph.* 30, 4 (July 2011), 27:1–27:10.
24. Obrenovic, v., and Martens, J.-B. Sketching interactive systems with sketchify. *ACM Trans. Comput.-Hum. Interact.* 18, 1 (May 2011), 4:1–4:38.
25. Read, J., MacFarlane, S., and Casey, C. Endurability, engagement and expectations: Measuring children's fun. *Interaction Design and Children* 2 (2002), 1–23.
26. Reilly, R. C. Is expertise a necessary precondition for creativity?: A case of four novice learning group facilitators. *Thinking Skills and Creativity* 3, 1 (2008), 59–76.
27. Rosin, P., and Collomosse, J. *Image and Video-Based Artistic Stylisation*. Computational imaging and vision. Springer London, Limited, London, United Kingdom, 2013.
28. Rubin, K., Fein, G., and Vandenberg, B. Play. In *Handbook of Child Psychology*, Wiley (New York, New York, USA, 1983), 693—774.
29. Ryokai, K., Marti, S., and Ishii, H. I/O Brush: Drawing with Everyday Objects As Ink. In *Proc. CHI '04* (2004), 303–310.
30. Shneiderman, B. Creativity support tools: accelerating discovery and innovation. *Communications of the ACM* 50, 12 (Dec. 2007), 20–32.
31. Sternberg, R. J. *Handbook of creativity*. Cambridge University Press, Cambridge, UK, 1999.
32. Takagi, S., Matsuda, N., Soga, M., Taki, H., Shima, T., and Yoshimoto, F. A learning support system for beginners in pencil drawing. In *Proc. GRAPHITE '03* (2003), 281–282.
33. Union, I. T. ITU-T Rec. P.910: Subjective video quality assessment methods for multimedia applications, 1996.
34. Union, I. T. ITU-T Rec. BT.500-11: Methodology for subjective assessment of the quality of television pictures, 2002.
35. Zeng, K., Zhao, M., Xiong, C., and Zhu, S.-C. From image parsing to painterly rendering. *ACM Trans. Graph.* 29, 1 (Dec. 2009), 2:1–2:11.

¹The first three images of Fig. 5 are from Flickr users “swainboat”, “Martino Pizzol”, and “Vinoth Chandar” (from left to right), they are released under a Creative Commons Attribution 2.0 licence (<https://creativecommons.org/licenses/by/2.0/legalcode>).