



RECODE: revision control for digital images

Fabio Calefato¹ · Giovanna Castellano¹  · Veronica Rossano¹

Received: 4 December 2018 / Revised: 15 March 2019 / Accepted: 6 May 2019 /
Published online: 16 May 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

Revision control is a vital component in the collaborative development of artifacts such as software code and multimedia. While revision control has been widely deployed for text files, very few attempts to control the versioning of binary files can be found in the literature. This can be inconvenient for multimedia applications that use a significant amount of binary data, such as images, videos, meshes, and animations. Existing strategies such as storing whole files for individual revisions or simple binary deltas, respectively consume significant storage and complex semantic information. To overcome these limitations, in this paper we present RECODE, a revision control system for digital images. It stores revisions in the form of a DAG (directed acyclic graph) where nodes represent editing operations, and edges describe the spatial and temporal relationships between operations. Being integrated with GitHub, the largest project hosting platform, RECODE also facilitates the artistic creation process of distributed teams with different workflows that include image editing and digital painting. A preliminary user study was performed to assess the perceived usability of the proposed system.

Keywords Multimedia design · Image editing · Revision control · Collaborative design · Digital painting

1 Introduction

In multimedia design and development, there is a wide range of contents such as text, images, video, and audio that need to be created and edited. Recently, collaborative forms of multimedia development have proved to be useful for authoring, editing, collecting, and producing digital content [7, 8, 20].

✉ Giovanna Castellano
giovanna.castellano@uniba.it

Fabio Calefato
fabio.calefato@uniba.it

Veronica Rossano
veronica.rossano@uniba.it

¹ Computer Science Department, University of Bari, Via Orabona 4, 70125, Bari, Italy

When the development of multimedia is carried out in collaborative and integrated design environments, *revision control* (or version control) becomes essential to calculate, represent, and store differences between successive versions of the developed digital objects. Indeed, the development process can involve many authors with potentially different skills. The current paradigm of collaborative editing relies on sharing files between stakeholders upon each revision. This requires maintaining consistency of the versions and dealing with concurrent edits in the same part of a multimedia artifact. Hence, the need for revision control in digital artifacts comes in two different forms: ① maintain a versioned history of changes applied to artifacts and ② enable the creation of several concurrent variants of the same artifact as well as their consolidation.

However, so far revision control has been widely deployed for text files, while visual media such as diagrams and pictures have received considerably less attention [17]. Even popular, commercial tools for image editing such as *Adobe Photoshop* offer little or no support for managing variation in the produced artifacts, forcing users to employ basic techniques to track multiple versions of their work, such as merely creating multiple copies of the picture file. Also, very few methods are available for the efficient storage of modifications of visual artifacts and in general of any binary content. Existing strategies such as storing whole files for individual revisions or simple binary deltas could consume significant storage and obscure semantic information. This can be inconvenient for multimedia applications that use a significant amount of binary data, such as images, videos, and animations.

In this work, we present *RECODE* (REvision CONTROL of Digital imagEs), a tool aimed to support image version control and collaborative creativity in digital multimedia projects. Inspired by version control systems for source code, such as *Git* and *Subversion*, where branching (i.e., creating a different version of a stored artifact) is a natural operation, we adopt the concept of paths and nodes in a graph to store persistent states over time owned by a particular image file. The core idea of the RECODE system is to store the history of editing operations applied to an image using a *Directed Acyclic Graph* (DAG) [1, 21], where each vertex contains the result of a performed editing operation. This representation is suitable for nonlinear revision control since the model has intrinsic support for collaborative editing, including branching and merging paths edited by multiple contributors. Also, the model does not require storing the whole modified image.

The remainder of the paper is organized as follows. In Section 2 we introduce the problem of revision control in visual artifacts as a background of our work. Section 3 overviews recent state-of-art works related to the proposed tool. In Section 4 we describe the proposed tool. Section 5 provides experimental results on usability tests. Concluding remarks are given in Section 6.

2 Background

Revision control systems are widely used tools in software development primarily aimed at managing versions of software source code during implementation [18]. In general, revision control tools assist the management of evolving digital artifacts, providing features that allow users to track intermediate revisions of artifacts and their relations. Revision control systems also facilitate distributed and simultaneous content creation through four basic operations: `diff`, `patch`, `branch`, and `merge`. The `diff` operation is used to extract the differences, or deltas, between any two versions of a given artifact. The `patch` operation is used to generate a specific version of an artifact based on a delta. The `branch` operation is




Fig. 1 An example of revision control with two parallel branches (A and B) created from a common ancestor and then merged

used to duplicate an artifact under revision so that modifications can happen in parallel and independently along two branches (Fig. 1). Finally, a *merge* is used to reconcile into one artifact two branches (i.e., parallel modifications) with a common ancestor (Fig. 1).

Depending on the adopted history model, that is, how the history of versions and changes to artifacts is stored, revision control systems can be classified as *state-based* or *change-based* systems [15]. Change-based revision control systems (also known as operation-based) store the operations actually performed between two succeeding versions in the revision control system. Instead, state-based systems store the history of changes as revisions of the versioned artifacts as they existed at various points in time. Most existing general-purpose revision systems such as Git¹ and Subversion (SVN)² employ a state-based model. These modern revision control tools save storage space by computing and persisting only the difference (i.e., delta) between succeeding revisions while preserving the full state of a few special versions – like initial or final (head) revisions [14].

Moreover, version control systems used for software development, such as Git and SVN are generally readily applicable to textual artifacts since source code is written in plain text format. However, these systems cannot handle the version control of binary files or offer limited built-in support for binary content [10]. When dealing with image data, in fact, they store separate images for each revision, thus wasting storage space. This issue alone hampers the adoption of revision control systems in managing digital images. Some extensions have been proposed (e.g., Git LFS),³ but they are meant to add support for storing large binary files rather than for versioning.

Besides, commonly used image-editing software tools offer minimal control for image versions. For example, *Adobe Photoshop* provides a version history that retains the undone operations on a stack. Hence, the user can quickly jump to any recent state of the image created during the current working session. However, high-level operations such as comparison (diff), branching, and merging of different versions are not available in Photoshop, nor other current drawing tools. A few other tools, such as *Adobe Drive* and *AutoDesk Vault*, offer version control for digital images, but they are proprietary and do not describe their specific model. Some more general version control tools, including *Perforce* and Git,

¹<https://git-scm.com>

²<http://subversion.tigris.org>

³<https://git-lfs.github.com>

offer support for image diff operations via services such as *GitHub*.⁴ However, even when image deltas or low-level image information such as pixel-wise differences are used, they still lack sufficient high-level semantic information for reviewing, branching, merging, or visualization.

3 Related work

Version control systems manage content change in a digital artifact and maintain a history of its evolution due to successive editing operations. Recovering editing operations for binary data is more difficult than for text, posing a big challenge for collaborative editing of digital visual artifacts. To address this, different revision control systems for visual artifacts have been recently proposed in the literature. The common idea is to implement a nonlinear control of versions using a DAG. Using a directed acyclic graph, the focus is on paths of editing operations rather than on pixels or image objects, and challenges such as differencing (or *diff* in short) and merge are solved with graph operations. They also offer support for selective undo and "nonlinear exploration", in which the user can adjust parameters to operations that have already been performed.

The first example of using a DAG for storing the history of operations (i.e., changes) and provide nonlinear revision control for binary image files is the seminal work by Chen et al. [6], who implemented a plug-in of the image editor GIMP,⁵ which tracks user editing actions in form of graphs to visualize revisions and support branching and merging.

A similar work is the one by Zhao et al. [23], who developed *skWiki*, a wiki-based framework for collaborative creativity in digital multimedia projects, including different types of media (text, hand-drawn sketches, and photographs). The framework uses the concept of paths as trajectories of persistent state over time. A document or file is represented as a path hence it is stored as the ordered (and timestamped) sequence of document-specific operations that created and modified it. Paths are implemented using a database management system (DBMS) that stores individual paths as tuples in one table, and all of the operations in another using the path identifier as a primary key and including the revision serial number.

The process of image differencing and merging concerns not only 2D images, but also 3D digital assets. Albeit not directly comparable to RECODE, a couple of studies are worth mentioning because of the use of DAGs. Doboš and Steed [10, 11] proposed a general approach to serialize DAG structures representing 3D assets and store them in MongoDB, arguing that NoSQL databases are better suited for storing spatial information. Our implementation is similar in that we also serialize DAGs into a database, albeit we opted for MySQL since we deal with 2D (i.e., non-spatial) assets only. Wang et al. [22] proposed a 3D scene editor that builds upon the approach proposed by Doboš and Steed. Similarly to our work, they also extended an open source web-based editor (three.js),⁶ store DAG structures into a database, and implement the typical features of a revision control system, including `commit`, `branch`, `diff`, and `merge`.

Instead of relying on DAG structures, da Silva et al. [9] implemented *IMUFF*, an image version control system that works at bit level – i.e., the delta between two images is computed as the set of bits that are changed across the two versions. Unlike DAG-based tools

⁴<https://github.com>

⁵<http://gimp.org>

⁶<https://threejs.org>

such as RECODE and skWiki, IMUFF has the advantage of being independent of the image editing tool. This, however, causes limitations too: ① a color image with a somewhat low resolution of 1024x1024 pixels requires over 4 million bitwise comparisons; hence, IMUFF has been designed to work only on architectures with high-end NVIDIA GPUs to take advantage of their massive parallel computational power; ② since there is no tracking of operations, as in RECODE and other DAG-based tools, IMUFF only supports the tracking of rigid transformations – i.e., differencing and merging works only for rotations, translations, and reflections – whereas tracking of filter effects such as blur and solarize is not supported.

Finally, besides academic prototypes, a few commercial solutions for image differencing and merging are also available. While primarily intended for source code revision, GitHub provides an image comparison tool that allows visualizing differences between stored images through a split pane view; the tool, however, does not support merging. *Araxis Merge*⁷ is a commercial tool that supports both text and image differencing and merge. Several image formats are supported and differences are shown at a pixel level, as in the case of the *IMUFF* prototype by da Silva et al. [9]. *Abstract*⁸ is a commercial tool that offers version control for files edited with Sketch,⁹ a proprietary vector graphics editor for macOS. With respect to the collaboration workflow, Abstract is very similar to RECODE. Both tools allow creating projects in a repository (proprietary in their case, GitHub in ours), add images assets (plus any type of files in RECODE), storing only important revisions as entire binary files (instead of deltas), and branching/merging alternative versions of images, even those created by other team members who have access to the same project repository.

Inspired by the work proposed in [6], in [5] we introduced the first prototype of a revision control system for digital images that implements a hybrid approach combining the use of graphs with both state-based and change-based revision control. Like in Chen et al. [6] and other change-based systems [15], we also expose the history as a DAG to represent spatial, temporal, and semantic dependencies between successive recorded image-editing operations that are stored as graph nodes. However, like state-based revision control systems, we also allow users to store select, important revisions as binary files, thanks to the integration with Git, so that users do not need to constantly reconstruct them by reapplying the entire sequence of operations that ultimately lead to their creation. The prototype presented in [5] was intended as a proof of concept for our hybrid model, hence it offered only a minimal set of graphical operations. In order to make our solution more robust, we developed a new version of our system that was obtained by integrating our hybrid model within a more sophisticated open-source image editor. As a result of this integration, in this paper, we present RECODE,¹⁰ a tool for nonlinear versioning control to be used in collaborative works involving digital image editing.

Although inspired by the work proposed in [6], RECODE differs from it in the following aspects:

- Chen et al. [6] developed their prototype on the GIMP image editing software, which, however, does not provide an API to access the history of operations; therefore, they had to fork the code and develop a custom version that is now obsolete – i.e., it has never been updated to be on par with the official version. To avoid a similar issue, we

⁷<https://www.araxis.com/merge>

⁸<https://www.goabstract.com>

⁹<https://www.sketchapp.com>

¹⁰<https://github.com/RECODE2/recode>

developed RECODE by extending *miniPaint*,¹¹ an open-source, online image editor that, while being not as sophisticated as GIMP and Photoshop, provides native support for extensions as well as access to the entire history of edit operations. As such, RECODE promises to be more future-proof.

- RECODE is integrated with Git and GitHub, thus it enables teams to collaborate in the creation of both visual and text artifacts. Besides, our system works well with the Git LFS extension, thus allowing users to check out only the entire image files needed for the current task at hand, and just symbolic references for other images that are not needed.
- RECODE can be seamlessly used online and offline. In fact, despite being a web application developed in JavaScript, we provide a Docker image that allows to execute the tool locally and to synchronize the local Git clone with the remote origin repository as soon as an Internet connection is again available.
- Instead of using custom formats as in [6], RECODE stores meta-data and project-related information using files in standard JSON objects, which are serialized into a MySQL database.

Finally, the RECODE tool satisfies the following fundamental properties required to software tools supporting collaboration around digital artifacts [23]:

1. **Mobility** – *platforms should be accessible from everywhere*; RECODE leverages a responsive user interface that adapts to the screen size of mobile devices such as tablets.
2. **Collaboration** – *platforms should allow collaboration between geographically and temporally distributed participants*; RECODE workflow can support collaboration between teams of distributed people who can work remotely and also online/offline by running the tool locally via a Docker container.
3. **Revision history** – *platforms should enable the tracking and versioning of digital assets*; RECODE supports the most important features available in revision control systems thanks to its tight integration with Git and GitHub.
4. **Transparency** – *platforms should hide revision systems complexity from users*; RECODE keeps complexity to a minimum by providing an intuitive UI to execute revision system commands such as `commit`, `branch`, and `merge`.
5. **Rich media** – *platforms should allow for the presence of different media types*; RECODE allows teammates with different interests (e.g., developers, illustrators, writers) to collaborate using the same Git repository, which can archive all types of artifact, whether textual or binary.
6. **Divergent creativity** – *platforms should enable the creation and tracking of different versions of the same artifacts*; RECODE supports divergent creativity by enabling branching in non-linear revision – i.e., representing alternative versions of an image as alternative paths in a DAG – as well as the differencing and merging of such ‘divergent’ artifacts.

4 The revision control system

The core data structure of RECODE is a Direct Acyclic Graph (DAG) that is used to store the revisions as deltas (Fig. 2). DAG nodes represent image editing operations with

¹¹<https://github.com/viliusle/miniPaint>

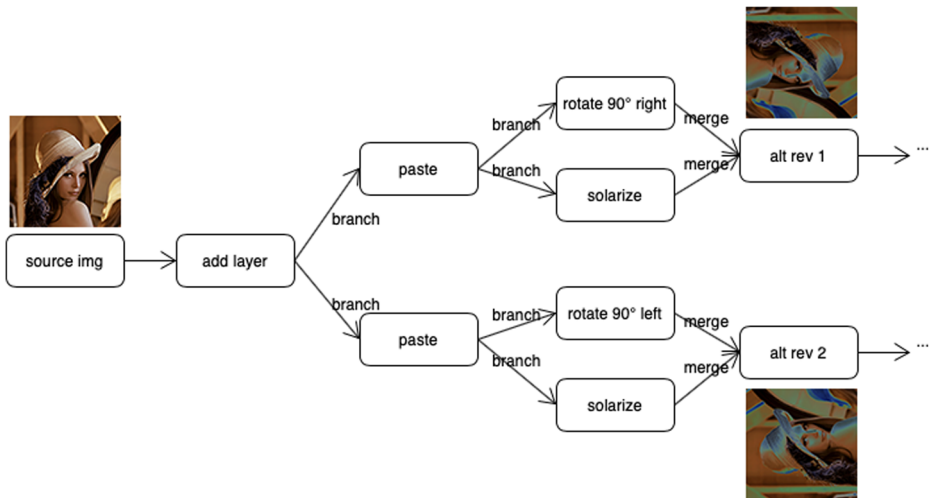


Fig. 2 An example of recorded editing operations stored as a DAG

relevant information such as the type of operation and its parameters, the author who applied the operation, the time of the application and eventual notes. DAG edges represent the relationships between the operations.

A (directed) sequential path between two nodes implies a spatial and/or semantic dependency between operations and the path direction gives information about their temporal order. Spatial dependency considers the spatial relationships between operations. Two operations are spatially independent if they are applied to non-overlapping regions. For example, *drawing a shape* and *coloring it* are spatially dependent operations. Conversely, *drawing a shape* and *coloring another existing shape* are independent operations. Semantically independent operations are rigid transformations (e.g., translation, rotation), deformation (e.g., scale, shear, perspective), color adjustment (e.g., hue, saturation, brightness, contrast, gamma) and filter (e.g., blur, sharpen).

Multiple parallel paths between two nodes imply independent operation sequences, namely those that apply on disjoint regions of the image. The DAG records the user editing operations and dynamically grows as more revisions are committed. Each revision in our system is a sub-graph of the DAG originating from the first node which represents the act of initialization – i.e., opening an empty canvas or loading an existing image. The state of the revision is always equivalent to the result generated by traversing its corresponding sub-graph. It should be noted that in our system, the DAG encodes only actions, not whole images. Table 1 lists the editing operations that lead to the creation of a new node in the DAG.

In the following, we briefly describe the main revision control commands implemented in RECODE.

4.1 Revision control commands

Based on the DAG data structure, RECODE provides the primary mechanisms for automatic resolving and merging multiple revisions with potential conflicts, as well as a user

Table 1 Some of the most common image-edition operations leading to the creation of a new node in the DAG

Type	Operation
Rigid transformation	Mirror, Flip, Transpose
Deformation	Scale
Color and filter	Histogram, Brightness, B&W, Sepia, Invert, Solarize, Posterize
Edit	Crop, Text, Reset
Brush	Brush
Load image	New, Import

interface that allows manual change and intervention on automatically merged images (see Section 4.2).

The implemented revision control commands include review, addition, branch, merge, difference, and conflict resolution. All these commands are offered through a friendly user interface.

4.1.1 Diff

While the classic line-based `diff` command [14] is commonly used to extract differences between text files, there is no such well-defined difference tool for images. Among general image comparison visualization approaches, popular ones include side-by-side comparison (e.g., *Adobe Bridge*, *Perforce*), layer-based difference (e.g. *PixelNovel*), per-pixel difference, image overlay (e.g., *Wet Paint* [3]), and flickering difference regions (e.g., the compare utility of *ImageMagick*). These approaches are designed to handle only low-level bitmap differences, with little information about the editing semantics.

In contrast, following the idea of Chen et al. [6], RECODE realizes an informative `diff` by leveraging all the relevant high-level information recorded in a DAG. As further illustrated in Section 4.2.1, all data structures are stored as JSON objects, including DAGs and image metadata. Therefore, to implement the `diff` operation between two image revisions, RECODE relies on the `fast-deep-equal` package¹² of Node.js to identify changes in all the fields of a couple of JSON files and stores the differences (i.e., the delta) in a new JSON file. This delta is designed not only to reduce storage space but also to allow for the fast identification of changes and conflict resolution during the integration of changes – i.e., the merge operation described next.

4.1.2 Merge

The merge operation is performed in version control systems to consolidate two artifact revisions created in parallel branches. Unlike text, merging two image versions (Fig. 3) requires complex procedures in order to identify the difference between them. In text files, changes are identified through line-by-line comparison. Instead, with binary files, it is difficult to define which parts of an image have been changed.

To implement the merge operation in RECODE, we relied on the `merge-json` package¹³ for Node.js. The package is used to merge the editing history of two different images

¹²<https://www.npmjs.com/package/fast-deep-equal>

¹³<https://www.npmjs.com/package/merge-json>



Fig. 3 An example of merging two versions of the same image

recorded into their corresponding JSON metadata (Fig. 4). Thus, in the resulting merged JSON file, all the operations are sequentially ordered in one branch. This is possible because all operations applied to an image are recorded using a growing integer as unique *id* and an attribute *order* to define the priority when changes involve an overlapping area of the image. Further information on the structure of JSON files are given in Section 4.2.1.

4.1.3 Git commands

The RECODE tool is developed to use Git repositories shared online on GitHub. When a new project is started, the URL of the associated GitHub repository must be provided, along with the path of a local folder where the content of the remote repository will be cloned locally. Then, it is possible to execute all the Git commands to control the image revisions as binary files both locally and remotely. The Git commands reviewed next are made available from the UI (see Section 4.2 for more).

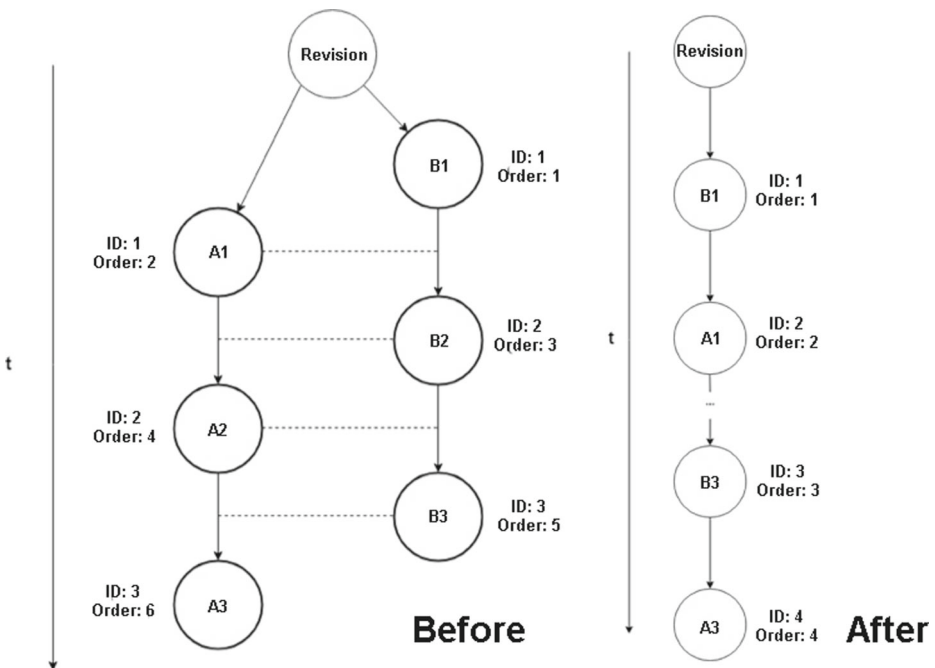


Fig. 4 RECODE: An example of a merge between image A and B (before), with operations and resulting revisions ordered sequentially in a single branch (after)

Since Git is a distributed revision system, the storing and sharing of artifacts are performed in separate steps; specifically, the `git commit` command is used to store files locally, whereas the `git push` command is used to send local changes to the shared remote repository, so that collaborators can retrieve them. Because this distinction may increase the cognitive load for the non-developer users, RECODE has been designed to simplify the workflow and hide away such complexity. Initially, when the repository is empty, the user adds the initial image and commits it as revision 0. To this end, RECODE provides a high-level ‘Add revision’ command from the revision control window, which locally stores the first image revision (i.e., the initial node of the DAG) as a binary file; the file is committed to the local Git repository by issuing under the hood a `git commit` command.

Committing revisions is one of the most frequently used revision control commands. To save the current work progress as further binary revisions, users can select the high-level command ‘Commit’ from the revision control window, which will save changes as a binary file both locally and remotely – i.e., by issuing first a `git commit` command, followed by a `git push`. Although users can save revisions whenever they like, it is generally unnecessary for them to do so in an action-wise fine-grained fashion, since RECODE can record all the actions and flexibly visualize them as a DAG. As a general guideline, users proceed to commit revisions when one of the following two conditions is met: 1) some milestone of the work is achieved or 2) users would like to try out different variations. In the latter case, the committed revision can be used as a branch point for future reference or revision rollback.

Furthermore, our system is compatible with the Git LFS (Large File Storage) extension that allows, when enabled, the storage of large files to a separate repository. The original repository will only contain pointers to the actual large files. Thus, a user can decide to download these files only if need be. As such, the use of Git LFS is recommended for speeding up the access to project repositories that host many very large files as in the case of images.

4.2 Architecture

RECODE is based on a client/server architecture that includes the following components (Fig. 5):

- the *Web server*, which executes the Node.js JavaScript runtime environment, where the miniPaint image editor is executed.
- the *Image Editor*, which extends miniPaint adding the components that enable the revision control workflow, and in particular:
 - the *Logger*, which silently records user-editing actions in the background in the form of JSON log files.
 - the *Replayer*, which replays actions in the image editor starting from the deltas stored in the log files;
- the *Database*, which stores all the information needed to create the revision graph and implement the revision control functions;
- the *Git and Git LFS Servers*, which physically store images files as binary artifacts and enable a distributed revision control.
- *Client browser*, which renders the tool UI and triggers the revision control commands.

In the following, we describe the custom components of RECODE in further details.




Fig. 5 Deployment diagram showing RECODE architecture and the distribution of its component to deployment targets

4.2.1 Image editor

To develop the image editor, we chose to extend an existing open source tool, *miniPaint*, which provides for an easy GUI and a wide set of editing tools, filters, and rigid transformations, which can be applied to images. Image editing is also based on layers and images can be saved in several common formats, including PNG, JPEG, BMP, and TIFF. *miniPaint* is a single page application that runs in Node.js, which is an I/O event-driven framework based on the Chrome JavaScript V8 engine. Despite Node.js is a server-side framework, *miniPaint* uses a specific module to compile the code in JavaScript and use it in a Web browser.

A distinguishing feature of *miniPaint* is that it can natively export all the sequence of actions made on an image into a JSON file structured as follows (Fig. 6):

1. *info* contains information such as the file name, the software release version, and the id of the last active layer.
2. *layers* is a list of all the layers existing in an image; for each layer the following setting info are recorded: id, the type of applied operation, the order (priority) of the layer with respect to other layers, a data field containing all points where the editing operation was applied, and the type field, which can be image when a new image is loaded or the type of the specific editing operation applied (e.g., brush).
3. *data* contains the numeric id of the image, as well as its MIME type and encoding in the data field (in the example, the image is in PNG format with a base64 encoding).

4.2.2 Logger and replayer

Two of the main components of RECODE developed to extend *miniPaint* are the logger, which records in text mode in the background all the actions that the user performed by a

```

{
  "info": {
    "width": 800,
    "height": 720,
    "layer_active": 2
  },
  "layers": [
    {
      "id": 1,
      "type": "brush",
      "opacity": 100,
      "order": 1,
      "data": [
        [
          99,
          101
        ]
      ],
      "color": "#5cabfb",
    },
    {
      "id": 2,
      "parent_id": 0,
      "name": "bicchiere.jpeg",
      "type": "image",
      "opacity": 100,
      "order": 2,
    }
  ],
  "data": [
    {
      "id": 2,
      "data": "data:image/png;base64"
    }
  ]
}

```

Fig. 6 An excerpt of the JSON file generated by miniPaint

user while editing an image, and the `replayer`, which allows the ‘replay’ of the actions stored in the DAG to visualize the final image in the editor. miniPaint easily support these two features since it can natively export into JSON format the sequence of actions applied to an image and is also able to portray an image in the editor by reading from its JSON file the actions performed by the user.

4.2.3 Database

A relational MySQL database (Fig. 7) has been adopted to ensure the persistence of meta-data, thus maintaining the semantic, spatial, and temporal mutual dependencies between artifacts, namely to store all the information needed to create the revision graph and to implement the revision control functions. Using the database, we track information about the logged users (via GitHub), the list of repositories created by them, the commits made by different users sharing the same repository, and the branches generated from each revision. As described earlier, RECODE stores meta-data and project-related information using files in standard JSON objects, which are then are serialized into the database.

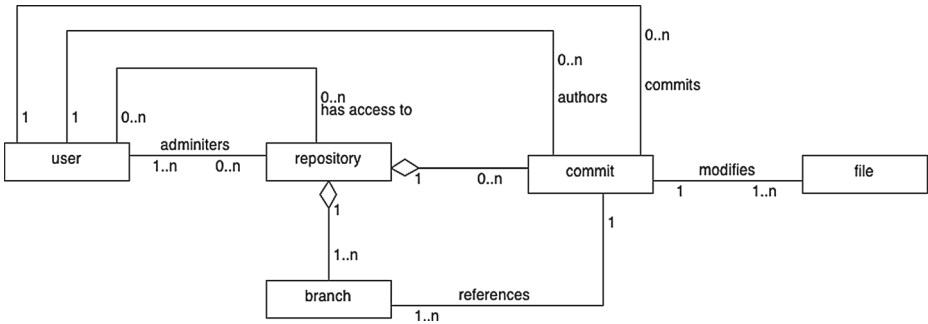


Fig. 7 A portion of RECODE conceptual schema to store the revision graph in the database

4.2.4 Git server

RECODE implements a hybrid approach combining the use of graphs with both state-based and change-based revision control. The DAG is used to represent spatial, temporal, and semantic dependencies between successive recorded image-editing operations that are stored as graph nodes. However, like state-based revision control systems, RECODE also allows users to store important revisions as binary files, thanks to the integration with Git, so that users do not need to continually reconstruct them by reapplying the entire sequence of operations that ultimately lead to their creation. This has been achieved by integrating RECODE with GitHub, a Git hosting provider, which allows sharing Git repositories. Thanks to the `git clone` operation, each user can download a copy of the repository from GitHub (via its API v3) and work locally and then commit back the changes.

4.3 Graphical user interface

The GUI of RECODE is based on the GUI of miniPaint. Position and layout of the menus are the same, with a few changes applied to accommodate the implementation of the extra features. Since miniPaint is not integrated with GitHub, we added support for user login and repository management. After authentication, users can access the revision control functions using the ‘Git’ item in the menu bar. The submenu includes ‘VCS’ and ‘Repository’ items.

The ‘VCS’ (Version Control System) item groups the following commands:

- *Revision Graph*, to visualize the DAG of a specific image (Fig. 8). The browser sends the JSON to the server for processing JSON in order to reduce the impact in terms of memory on the client side. As explained earlier, only the differences between two consecutive revisions are stored in a JSON file. Accordingly, when the Revision Graph of an image is loaded, all the differences along the path to the initial revision (i.e., the yellow node in Fig. 8) are concatenated to create the resulting JSON representation of the image to be displayed. The DAG viewer is a key feature of RECODE, which offers a visual management interface for navigating the evolving graph structure, with alternative paths generated during the editing of different image revisions.
- *Add Revision*, to commit an initial image revision and start a new workflow (Fig. 9);
- *Commit*, to store the binary version of an image and share it with the other members of the team via GitHub. Each commit requires a mandatory textual description to allow collaborators to understand what changes has been made to the image.

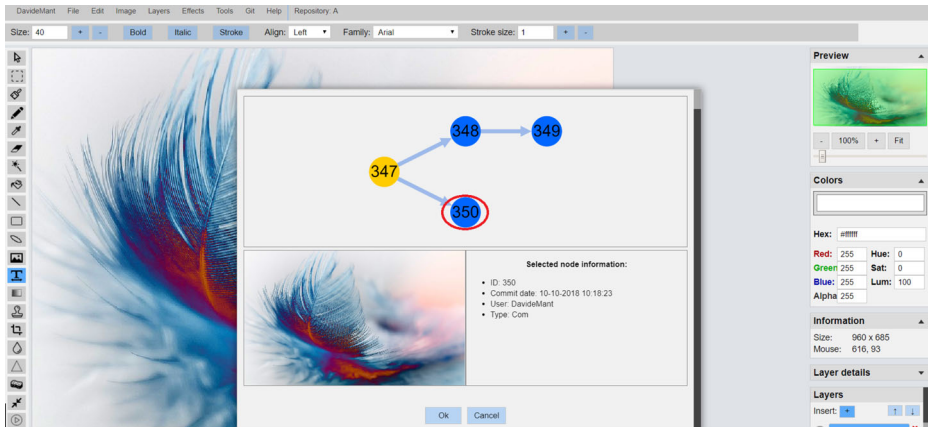


Fig. 8 GUI for the DAG viewer opened after a *Revision Graph* command

- *Merge*, to combine different image revisions. This function is used to combine changes from either the same user or other teammates who have modified the same image. This is graphically represented as merging two branches of the same graph by creating a new node that contains the sum of all contributions, as depicted in Fig. 10). Conflicts are solved by assigning a priority to each node in the DAG.

The ‘Repository’ menu item groups the commands related to the management of the GitHub repository in which all the images will be stored. The available functionalities are:

- *Create*, adds a new repository in GitHub;
- *List/Choose*, shows the list of user’s repositories to choose which to work on;
- *Clone Git repo*, creates a local copy of the working Git repository;
- *Participating users*, lists the users that can contribute to the image editing process;
- *Invite user*, allows the owner of the working Git repository to invite other collaborators;
- *Delete user*, allows the owner of the repository to delete one or more participants.

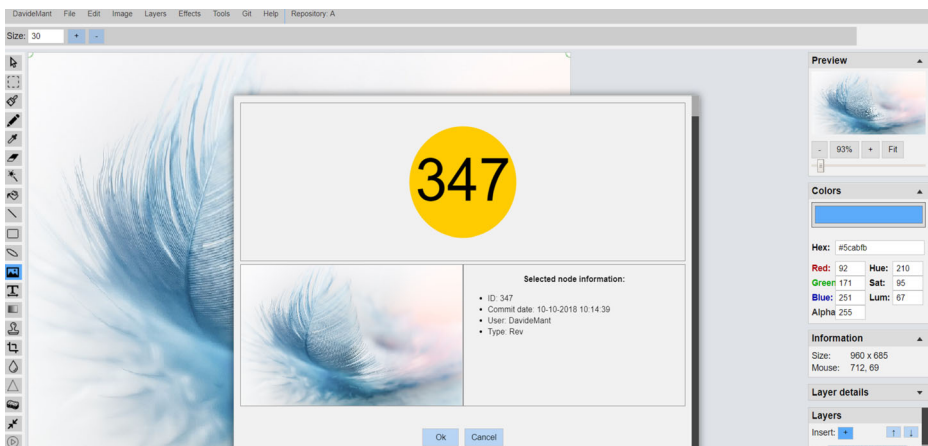


Fig. 9 GUI for the *Add Revision* command

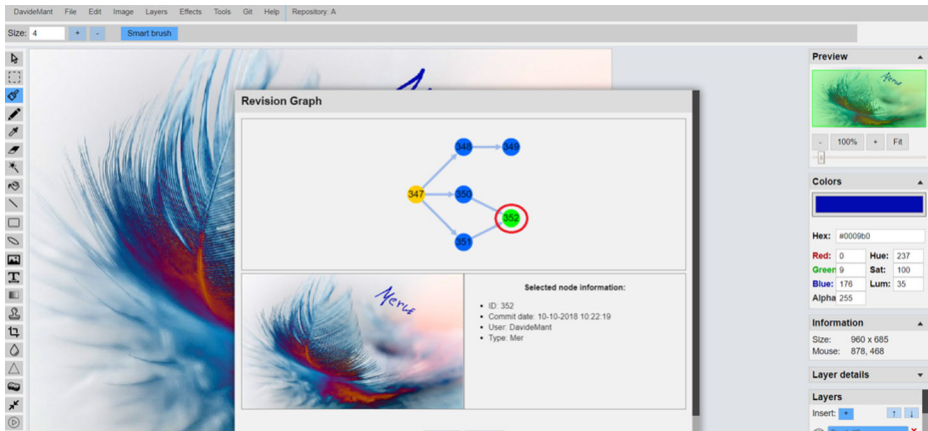


Fig. 10 Visualization of the DAG after the *Merge* command

5 User test

In order to evaluate the system, a usability test has been conducted. The adopted procedure is compliant with the eGLU usability protocol used by the Italian Public Administration [13]. This protocol, which is one of the most used for simplified usability tests, has been enriched by the eGLU-M for mobile systems and by the Usability Glossary of WikiPA project [13]. The protocol has been defined in order to be adaptable to different kinds of software. In this case, the following documents related to eGLU 2.1 protocol have been used:

- data of participants;
- description of the task;
- questionnaire for computing the Net Promoter Score (NPS);
- questionnaire for computing the System Usability Scale (SUS);
- table of results.

5.1 Study participants

The participants involved in the study were 15 subjects interested in image editing. In particular, the participants were 3 graphic designers who usually use professional tools for image editing, 9 computer scientists who are familiar with revision control systems, and 3 users interested in using image editors but without any specific background. As argued by Nielsen [16], 5 users are expected to find 85% of the usability problems of a system, hence a sample of 15 users can be considered a safe choice to collect significant results.

5.2 Assigned tasks

Each user was asked to complete the following tasks:

1. Create an account, create a project and a repository;
2. Open an image and commit the first Add Revision.
3. Apply three changes to the image and commit the result;
4. Starting from an image revision (a node in the DAG) create a new branch, apply some changes, and commit the result;

5. Visualize the information about a commit and load an image in the editor;
6. Apply the `merge` operator to unify the changes between two branches;
7. Invite a new participant to the repository;
8. Load an image, apply a change, and create a new revision.

We observed that the interaction with the system was not easy for all users. Some of them encountered difficulties using both the image editor and the version control commands. This aspect was detected through both the conductor's observation, who supported the users during the test, and quantitative measures, such as the time spent (in seconds) by users to perform each task. Figure 11 shows the min, max, and mean execution times, aggregated per task.

The most time-consuming tasks were Task#1 and Task#4. Task#1 required the creation of a GitHub account, which involved several steps, from filling out the registration form to activating the account via email address verification. Users who happened to have a GitHub account already took considerably less time since they only had to log in. Task#4 required the users to create a revision and then add a new branch. We noticed that this actions required users to have specific knowledge in how a revision system works therefore subjects without a computer science background struggled to complete this task. In other cases, such as Task#3 and Task#8, some users lingered before applying the requested changes to the image since they chose to fiddle with the tool trying to create a nicer image.

Overall, the analysis of the results highlights that those users who were not already familiar with version control systems spent more time than expected.

5.3 Net promoter score

After carrying out the tasks, the Net Promoter Score (NPS) was used to measure the users' overall usefulness perception of the system. The NPS is calculated by collecting the answers to the question: "*How likely is it that you would recommend RECODE system to*




Fig. 11 Average time spent by participants on each task




Fig. 12 SUS score obtained from the usability test on the 15 participants (P1,P2,...,P15). The line indicates the average SUS score

a friend or colleague?” The answers were calculated using a 0-10 scale. Respondents can be *promoters* (score 9-10), i.e., enthusiasts of the system, *passives* (score 7-8), i.e., satisfied but unenthusiastic users, *detractors* (score 0-6), i.e., unhappy users.

The results show that 33.33% of the participants belong to promoters, 53.33% are passives and the 13.33% are detractors. Subtracting the percentage of *detractors* from the percentage of *promoters* yields the NPS, which can range from a low of -100 (if every customer is a *detractor*) to a high of 100 (if every customer is a *promoter*). In our case, the resulting NPS is 20, which is rated good according to existing literature [12].

5.4 System usability scale

The System Usability Scale (SUS) is a quick and reliable questionnaire for measuring the usability of an artifact (software code or multimedia) [4] and is currently used to evaluate the usability of different kinds of products and services. The questionnaire consists of 10 standard items,¹⁴ using a 5-point Likert scale.

Scores given by the participants are processed in order to derive a single value using the following procedure. Let x and y denote the score given by the user for odd items (1, 3, 5, 7, 9) and even items (2, 4, 6, 8, 10) respectively. The SUS score z is computed using the following formula:

$$z = 2.5 \cdot [(x - 1) + (5 - y)]$$

The resulting value z ranges in [0, 100] and represents the average satisfaction level of the user.

Based on existing literature [2], a SUS score above 68 can be considered above the average. In our case, an average SUS score of 68.9 was achieved (Fig. 12). This indicates that the RECODE tool has a quite good usability level even though some improvement

¹⁴<https://www.measuringux.com/SUS.pdf>

would be welcome. Actually, low scores were mainly obtained for the question “*I needed to learn a lot of things before I could get going with this system.*”. As a matter of fact, some users encountered difficulties in understanding terms related to collaborative tools (commit, repository, and so on), others have difficulties with terms related to the DAG (add a branch). Moreover, during the test some users were confused by similar terms used in the interface, for example layers and levels, and merge and merge down. Changes in the naming of RECODE functions, could fix these problems and improve the usability level of the tool.

6 Conclusions

In this work, we have presented RECODE, a tool for the revision control of digital images. Unlike other existing systems for image revision control, RECODE adopts a hybrid approach that saves user editing actions as direct acyclic graphs (to save storage space) but also allows users to save important milestones revisions as binary files. Thanks to the integration with the popular revision control system Git, our system is capable of supporting both collaboration and revision control. Hence, we argue that the tool can be adopted in distributed collaborative team works by virtually any audience involved in digital content creation, from web designers to engineers, researchers, and creative artists. The usability study conducted with a few subjects provided initial evidence that our revision system has some usability problems. A possible cause can be found in the composition of the sample, since all the participants were not expert in collaborative activities. Moreover, we underline that the aim of this work was to develop a novel revision control system for multimedia, to fill a gap in the state of the practice. Hence, our work has focused so far more on the development of the tool that fulfilled its purpose, devoting less attention to other aspects such as the usability. Further work is currently in progress to improve the usability of the system, and in particular to adapt the interface terminology, which was one of the main difficulties reported by the users during the test session. To address this issue, we intend to apply the Participatory Design Approach [19], so as to involve the final users, namely graphic designers and multimedia developers. In this way, we expect to make more intuitive the usage of the tool, thus increasing the user satisfaction. Moreover, further tests have been planned to evaluate our system by involving users in real collaborative scenarios of multimedia development.

Acknowledgements This work is partially funded by the project “Creative Cultural Collaboration” (C3) under the Apulian INNONETWORK programme, Italy.

References

1. Bang-Jensen J, Gutin GZ (2008) Digraphs: Theory, Algorithms and Applications. Springer Science & Business Media, Berlin
2. Bangor A, Kortum PT, Miller JT (2008) An empirical evaluation of the system usability scale. *Int J Hum Comput Interact* 24(6):574–594
3. Bonanni L, Xiao X, Hockenberry M, Subramani P, Ishii H, Seracini M, Schulze J (2009) Wetpaint: scraping through multi-layered images. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, pp 571–574
4. Brooke J et al (1996) Sus-a quick and dirty usability scale. *Usability Evaluation in Industry* 189(194):4–7
5. Calefato F, Castellano G, Rossano V (2018) A revision control system for image editing in collaborative multimedia design. In: *Proceedings of the 22nd International Conference on Information Visualisation (IV 2018)*, pp 512–517

6. Chen HT, Wei LY, Chang CF (2011) Nonlinear revision control for images. In: ACM Transactions on Graphics (TOG), vol 30:4. ACM, p 105
7. Chen CW, Peng JW, Kuo CM, Hu MC, Tseng YC (2018) Ontlus: 3d content collaborative creation via virtual reality. In: International Conference on Multimedia Modeling, Springer, pp 386–389
8. Claman TH, Coniglio SJ, Daigle S, Gonsalves RA, Wallace RC (2018) Methods and systems for collaborative media creation, us patent no 9864973
9. da Silva Junior JR, Clua E, Murta L (2016) Efficient image-aware version control systems using gpu. *Software: Practice and Experience* 46(8):1011–1033
10. Doboš J, Steed A (2012) 3d revision control framework. In: Proceedings of the 17th International Conference on 3D Web Technology, ACM, pp 121–129
11. Doboš J, Steed A (2012) Revision control framework for 3d assets. *Eurographics*
12. Grisaffe DB (2007) Questions about the ultimate question: conceptual considerations in evaluating reichheld's net promoter score (nps). *Journal of Consumer Satisfaction, Dissatisfaction and Complaining Behavior* 20:36
13. Gruppo di lavoro per l'usabilità (glu), linee guida di design per i servizi web della pa (2018). <http://www.funzionepubblica.gov.it/glu>
14. Hunt JJ, Tichy WF (1998) Addendum to delta algorithms: an empirical analysis. *ACM Trans Softw Eng Methodol (TOSEM)* 7(4):449
15. Kleine M, Hirschfeld R, Bracha G (2012) An abstraction for version control systems. Universitätsverlag Potsdam
16. Nielsen J (2012) How many test users in a usability study. Nielsen Norman Group 4(06)
17. O'Sullivan B (2009) Making sense of revision-control systems. *Queue* 7(7):30
18. Ruparelia NB (2010) The history of version control. *ACM SIGSOFT Software Engineering Notes* 35(1):5–9
19. Schuler D, Namioka A (1993) *Participatory Design: Principles and Practices*. CRC Press, Boca Raton
20. Slaughter DS, Murtaugh MC (2018) Collaborative management of the elearning design and development process. In: *Leading and Managing e-learning*, Springer, pp 253–269
21. Thulasiraman K, Swamy M (1992) Acyclic directed graphs. *Graphs: Theory and Algorithms* 118
22. Wang Z, Cai H, Bu F (2014) Nonlinear revision control for web-based 3d scene editor. In: 2014 International Conference on Virtual Reality and Visualization (ICVRV), IEEE, pp 73–80
23. Zhao Z, Badam SK, Chandrasegaran S, Park DG, Elmqvist NL, Kisselburgh L, Ramani K (2014) skwiki: a multimedia sketching system for collaborative creativity. In: Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems, ACM, pp 1235–1244

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Fabio Calefato is an assistant professor of computer science at the University of Bari, Italy, holding the national scientific qualification as an associate professor. He is the General Chair of ICGSE 2019. He has been the co-organizer of the “Trust in Virtual Teams” workshop co-located with CSCW’13 and SSE’16 (Social Software Engineering) workshop co-located with FSE’16. He has served on the PCs of many software engineering conferences, including ICSE, XP, and ESEM. He also serves as a reviewer for the most important journals in the field of software engineering, such as TSE, EMSE, JSS, IST, and JSEP.



Giovanna Castellano is Associate Professor in Computer Science. She is the coordinator of the CILAB (Computational Intelligence Lab) at the Computer Science Department of the University of Bari. Her research interests are in the area of Computational Intelligence and include Fuzzy systems, Fuzzy Image Processing, Fuzzy clustering, Image processing and retrieval. She is the coauthor of the book “Fuzzy Logic for Image Processing: a Gentle Introduction using Java” (Springer, ISBN 978-3-319-44130-6) and co-editor of the book ‘Web Personalization in Intelligent Environments’ (Springer, ISBN 978-3-642-02793-2). She is Associate Editor of the journals: Information Science and International Journal of Systems, Control and Communications. She is on the editorial board of Journal of Knowledge-Based and Intelligent Engineering Systems and International Journal of Knowledge and Web Intelligence. She has served as a PC member for several refereed conferences or workshops in the field of Computational Intelligence, such as FUZZ-IEEE, ICANN, IJCCI, WILF, FCTA.



Veronica Rossano is Assistant Professor at Department of Computer Science at the University of Bari. Her research activities are focused on Educational Technology, in particular, the main aim of her research is to advance the state-of-the-art and practical use of computation and communication technologies for learning and teaching defining and validating new methods and techniques for the design and implementation of innovative learning environments. She is part of editorial staff of Journal of e-Learning and Knowledge Society and co-organizer of the International Symposium on “Multimedia & E-Learning” 2019 of IV2019. She has been program chair of different international and national conferences and workshops in the Educational Technology field. She has served as PC member of many conferences, including ICALT, ICEIT, and T4E . He also serves as a reviewer for the most important journals in the field of Educational Technology, such as Computers and Education, the Journal of Distance Education Technologies, the Journal of Visual Languages and Computing, and IEEE Multimedia.