Proposal for evaluating the effectiveness of semaphoric-tangible and virtual reality interfaces for a lo-fi previsualization activity

JOSHUA RAMSBOTTOM, UNIVERSITY OF CAPE TOWN
KATHERINE RIX, UNIVERSITY OF CAPE TOWN
DAVID RIX, UNIVERSITY OF CAPE TOWN

• Human-computer interaction ~ Multimodal user interface • Human-computer interaction ~ Virtual reality

Additional Key Words and Phrases: tangible, semaphoric, gestural, model-surface registration, iterative closest point (ICP), fully-occluding head-mounted display (HMD), 3D printing, Kinect for Xbox One (K4X1), Myo armband, Oculus Rift DK2, Blender, Maya

ACM Reference Format:

Joshua Ramsbottom, Katherine Rix, and David Rix. 2015. *Proposal for evaluating the effectiveness of semaphoric-tangible and virtual reality interfaces for a previsualization activity.* Honours report, University of Cape Town. 28 pages. June 2015.

1. DESCRIPTION AND CONTEXT

Previsualization (previz) is the process of generating an video mockup of the final film, based on an existing storyboard. The purpose of the previsualization stage is to use the mockup to review scene blocking, camera framing and pacing. Although the final output is a traditional 2D video, the 3D scenes used to generate the video must be editable, so that a variety of camera angles and scene layouts can be explored early in the creative process. (Changes in the production phase are more expensive.)

Previz output may be categorised as lo-fi or hi-fi. [Lwabona 2015] Lo-fi previz is characterised by a lack of animation detail. This could include: no attention to lighting; pose-to-pose cuts rather than interpolated animation; no limb articulation (characters move through the scene but remain in the “T” pose). Hi-fi previz does include animation details (such as articulated, animated limbs) but surface detail and other production-quality rendering is still ignored.

In some cases a mixed-fi approach is taken, varying the previz fidelity based on the demands of a particular scene (for example, the “blocked animation” (previz) sequence for *Big Buck Bunny* seen in [Goedegebure 2008]).

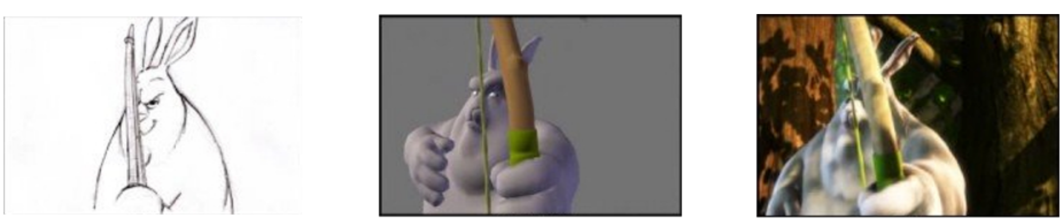

Figure 1.1. From left to right: storyboard image, previz frame, and production (final) frame in *Big Buck Bunny* (2008). Note how the visible detail increases at each stage.

2. PROBLEM STATEMENT

Animated films like *Toy Story* (1995) and *Inside Out* (2015) use digitally rendered 3D assets. The same complex software packages used in production are used to generate previz in pre-production. These software packages (such as *Blender* and *Maya*) must be operated by trained animators, who could otherwise be dedicated to production activities. They typically employ deskbound, keyboard-and-mouse, WIMP-based [1] interfaces.

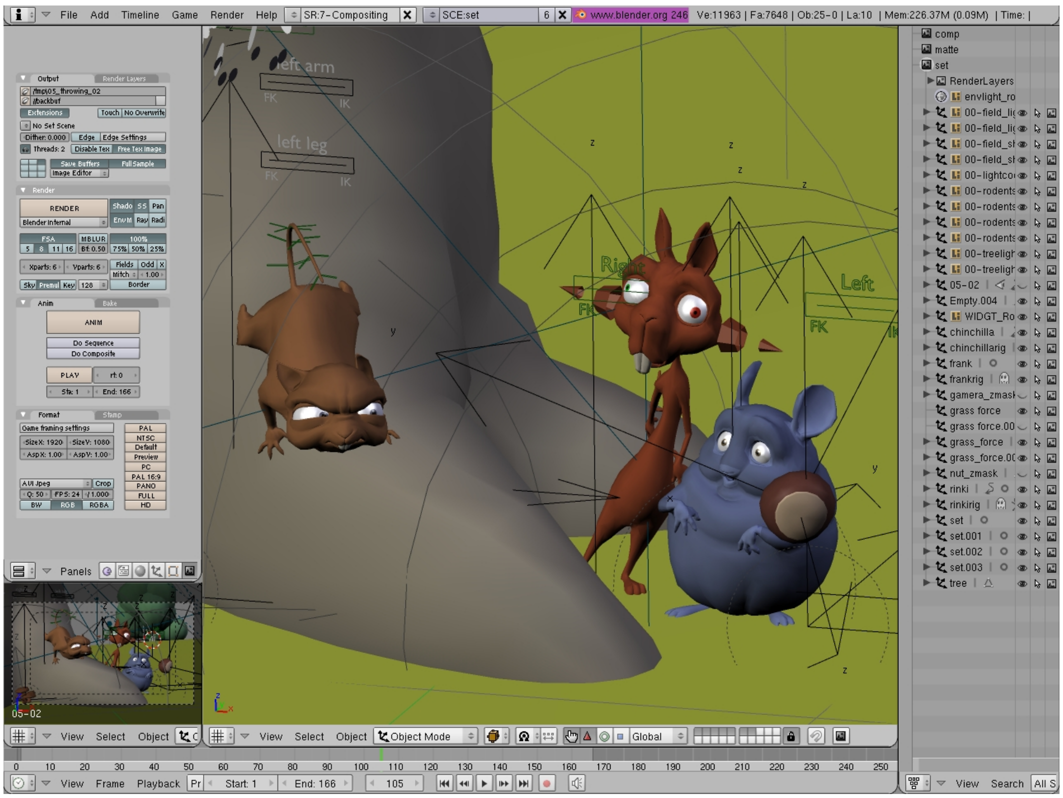

Figure 2.1. Blender application in use. Note the WIMP-based paradigm and the complexity of the interface (which requires training to use).

An interface that would allow non-animators with film expertise (for example, directors and producers) to effectively produce low-fidelity animation is therefore desirable.

More generally, there is a trend for non-professionals to produce creative work in previously inaccessible media [Tanenbaum et al. 2013]. Thus there is a general motivation to research interfaces suitable for a broader userbase.

1 WIMP refers to the familiar 'window, icon, menu, pointing device' paradigm found in the GUIs of most desktop applications. [Jacob et al. 2008]

This trend is enabled in part by the increased availability of commodity hardware, making previously exotic interfaces more feasible in the home and office setting, and more feasible for small-scale research projects [Berger 2013].

2.1 Alternate interfaces

We propose two systems that aim to leverage the existing 3D and reality-based skills [Jacob et al. 2008] of an untrained or casual user:

2.1.1. Tabletop system. The tabletop system makes use of solid character models, a wall-mounted GUI and gestural input. The user arranges the models on a tabletop, and performs semaphoric [2] gestures to control other aspects of the system. The position and orientation of the models is captured with a depth sensor, and displayed in the GUI on a large, wall-mounted screen.

Once the previz activity is complete, a mockup video is generated, which the user or a third party can review immediately. Any adjustments can be made immediately and easily.

A physically manipulated interface like this one is known as a tangible user interface (TUI). A system that combines modalities is known as a multimodal user interface (MMUI). We refer to this particular combination as “semaphoric-tangible”.

A more detailed introduction to this system is presented as a visual narrative in *Appendix B: Tabletop previz: how does it work?*

2.1.2. Virtual reality (VR) system. The virtual reality system places the user inside the scene with the virtual 3D assets. A head-mounted display (HMD) with a head-tracking subsystem is worn by the user, allowing them to look around inside the scene. They use a controller device to: move through the scene, move and interact with the assets (including the scene camera) and perform any other previz activities.

As with the tabletop system, a mockup video can be generated, reviewed and acted upon immediately and easily.

2.2 Research aims

Although the underlying interfaces are supported by the literature [Butterworth et al. 1992; Ishii 2008; Oviatt 1999; Oviatt et al. 2004], real-world applications of TUI and VR systems are uncommon and not always successful [Norman 2010]. In addition, animation activities must support domain-specific tasks (for example, judging how well a scene is framed).

Thus our aims are to make informed design decisions before implementing the proposed systems, and to measure their effectiveness through a repeatable user evaluation experiment. This experiment will compare the new systems to a pre-existing WIMP-based system. To this end we have:

---

2 Semaphoric input allows users to interact with a system by making particular gestures with their hands or bodies. The use of symbolic gestures distinguishes it from other gestural modalities, such as deitic input (pointing) or direct manipulation. [O'Hara et al. 2013]

- identified key tasks for the lo-fi previz activity (see *2.3 Key tasks*)
- formulated hypotheses that evaluate these interfaces (see *2.4 Research questions*)
- outlined a design process that incorporates domain expert and heuristic evaluation (see *3. Procedures and methods*)
- outlined a user evaluation experiment that tests our hypotheses, highlighting the what will be required to make our experiment repeatable (see *4. Evaluation* and *5.2 Repeatability*)

2.3 Key tasks

To create a previz mockup, a sequence of *keyframes* is created. Each keyframe is defined by a virtual scene arrangement, as seen from a virtual camera. In sequence, these keyframes tell the same narrative as the original storyboard. Additionally, the keyframes are interpolated to generate the final animated video. The choice and arrangement of keyframes is critical in capturing narrative, framing shots and timing the scene [Lwabona 2015].

In a previous review of the previsualization activity, we have identified a core subset of tasks that a lo-fi previsualization system should support:

2.3.1. Model placement is the task of placing and orienting characters (models) in a 3D scene in a single keyframe.

2.3.2. Camera placement is the task of placing and orienting the virtual camera in a 3D scene in a single keyframe.

2.3.3. Keyframe CRUD encompasses the tasks of creating, reading, updating and deleting individual keyframes in a sequence of keyframes. This includes modifying keyframe attributes such as timing.

2.3.4. Timeline navigation is the task of navigating backwards and forwards through the keyframe sequence. This also includes *scrubbing* (repeatedly previewing a short subsequence) and *playback* (previewing the entire sequence).

2.4 Research questions

With these aims and key tasks in mind, we will construct a scenario where a casual user (with no film expertise or training in animation software packages) must accurately translate a storyboard into a previz mockup video. In the context of this scenario, we have formulated the following research questions and hypotheses:

2.4.1. Tabletop interface. Is a semaphoric-tangible interface more usable than a WIMP interface when performing previsualization tasks?

Here the semaphoric-tangible interface is the one described in *2.1.1 Tabletop system*, using 3D-printed models, a Kinect for Xbox One (K4X1) depth sensor and a Myo electromyographic armband for semaphoric input.

To measure usability, we will test the number of errors and reported ease-of-use of the system (including user fatigue). We will also measure the efficiency of the system, by measuring the time taken to complete each subtask. We expect that the tabletop system to outperform the WIMP system on usability, but that the WIMP system will be more efficient.

HYPOTHESIS 2.4.1.1. The semaphoric-tangible interface will generate fewer user errors than the WIMP interface.

HYPOTHESIS 2.4.1.2. Users will report that the semaphoric-tangible interface is more usable than, and preferable to, the WIMP interface.

HYPOTHESIS 2.4.1.3. All subtasks will take less time when using the WIMP interface.

HYPOTHESIS 2.4.1.4. Expert evaluation will confirm that there is no significant difference in the accuracy of previz videos created using the semaphoric-tangible interface versus those created using the WIMP interface.

2.4.2. Tabletop depth sensor. The depth sensor is part of a critical subsystem in the tabletop interface and requires independent evaluation. Here we ask, can a K4X1 sensor perform real-time registration of multiple rigid but movable models in a dynamic tabletop environment?

In our implementation, the sensor will be operating at a distance of approximately 5m, and be required to register 3D printed models between 5 and 10cm tall. The models will be cartoon-like and have sufficient fine detail to appear in a 3D animated short. When the system is in use, there will be one human operator moving around the table, occasionally occluding the models.

HYPOTHESIS 2.4.2.1. At a range of 5-6m, the Kinect for Xbox One depth sensor will be able to correctly register a single 3D printed model.

HYPOTHESIS 2.4.2.2. At a range of 5-6m, the Kinect for Xbox One depth sensor will be able to correctly register multiple 3D printed models in a single scene.

HYPOTHESIS 2.4.2.3. At a range of 5-6m, the Kinect for Xbox One depth sensor will be able to correctly register multiple 3D printed models in a single scene, including re-registration after a human operator moves and/or occludes one or more models.

2.4.3. VR interface. Can a head-mounted virtual reality display-based system with the best input device create an effective and easy-to-use 3D user interface for previsualization?

The VR display used in this system will be the Oculus Rift, a fully occluding head-mounted display. There are a number of options regarding the input device to be used and thus more research will have to be done to determine the best input device for this system. Some options include an Xbox 360 controller, a Leap Motion gesture-based device, and a touch-based solution using a smartphone with motion tracking sensors. The effectiveness of the system will be measured using task performance, which includes both task completion time and error rates.

HYPOTHESIS 2.4.3.1. The VR interface will cause less errors than the WIMP interface.

HYPOTHESIS 2.4.3.2. Users will score the VR interface higher the WIMP interface when evaluating ease-of-use.

HYPOTHESIS 2.4.3.3. All subtasks will take less time when using the WIMP interface.

HYPOTHESIS 2.4.3.4. Expert evaluation will confirm that there is no significant difference in the accuracy of previz videos created using the VR interface versus those created using the WIMP interface.

2.4.4. Comparison. How do the tabletop and VR systems compare to a traditional mouse/keyboard-based WIMP approach?

Comparisons will be made with regards to both effectiveness and ease-of-use. Effectiveness in this case is defined as user task performance, which includes both completion time and error rates. Comparisons will first be made between each new system and the traditional one, and then between the new systems. We believe that both systems will improve upon the traditional system both in terms of effectiveness and ease-of-use.

HYPOTHESIS 2.4.4.1. Each alternate interface will generate fewer user errors than the WIMP interface.

HYPOTHESIS 2.4.4.2. Users will report that each alternate interface is more usable than, and preferable to, the WIMP interface.

HYPOTHESIS 2.4.4.3. All subtasks will take less time when using the WIMP interface.

HYPOTHESIS 2.4.4.4. Expert evaluation will confirm that there is no significant difference in the accuracy of previz videos created using the each alternate interface versus those created using the WIMP interface.

3. PROCEDURES AND METHODS

3.1 Design features

The new systems should both allow a non-animator presented with a storyboard to create a roughly animated sequence informed by that storyboard. They should support all of the key tasks already identified, including the arrangement of 3D scenes, manipulation of keyframes (including adjusting the timing of keyframes), and timeline navigation. The new systems should be usable with minimal training.

3.2 Development platform

We have selected the Unity 5 game engine, running on the Windows 8.1 operating system as a development platform. This choice was made based on the core functionality offered by the game engine (3D scene navigation and manipulation), existing skillset of the development team, and compatibility with the selected hardware devices (the Oculus Rift DK2 requires Windows 7 or higher, the K4X1 native SDK requires Windows 8 or higher, all devices have library support in Unity 5).

3.3 Implementation strategy and challenges

We will need base classes that can store the scene layout for a single keyframe, store a sequence of keyframes, and interpolate between keyframe scenes. These classes should be created such that they can be reused in each system (tabletop and VR), supporting the key tasks described above.

```
[       TABLETOP SYSTEM        ] [       VR ENVIRONMENT       ]
[ SEM. INPUT ] [ REGISTRATION ] [ 3D NAVIG. ] [ 3D/2D INPUT ]
[                  BASE CLASSES (core tasks)                  ]
```

Figure 3.3.1. Component diagram showing the relationship between the final systems, subsystems and base classes.

After implementing these classes, each system will be developed independently, but following the same process:

- We have already performed preliminary task analysis and system evaluation through expert interview, paper prototyping and limited interactive prototyping.
- We will perform further user requirement investigation through on-site training at an animation studio: Triggerfish Animation Studios, creators of *Adventures in Zambezia* (2012) and *Khumba* (2013).
- There will be 3-4 development iterations informed by the collected user requirements.
- After each iteration we will seek expert heuristic analysis. This feedback will allow us to refine interaction design as part of the development process.

- Finally, effectiveness of the completed systems will be measured via user evaluation.

By including limited user-centered design techniques in the development process, we will ensure that the system functionality is meaningfully chosen. Heuristic evaluations will guide the usability of the final systems. The final evaluation will measure the effectiveness of the interfaces for the storyboard activity.

Our primary challenge will be implementing these alternate interfaces with recently released devices. While they have active support channels and development communities, they are relatively new and may have undiscovered limitations and flaws. Additional challenges are listed under *8.4 Risks*.

4. EVALUATION

4.1 Registration accuracy

The registration subsystem requires separate evaluation before it can be incorporated into the tabletop system. The K4X1 sensor will be presented with 3D printed models on a tabletop. The calculated registration of models will be compared to their measured position and orientation. Registration will take place in three contexts: single model on a tabletop (repeated for all models), multiple models on a tabletop at the same time, and multiple models on a tabletop with a human user intervening, modifying the scene, and occasionally standing in front of the depth sensor.

Each scenario will be repeated 40 times, and the results evaluated for statistical significance.

4.2 System effectiveness

In the system evaluation, 80 participants will be presented with the same activity: the creation of a previz sequence based on a predetermined storyboard. Half the users (40 participants) will use and compare the tabletop and WIMP systems, the other half the VR and WIMP systems. To account for learning effects, half of each subpopulation (20 participants) will evaluate the systems in alternate-then-WIMP order, the other half in WIMP-then-alternate order.

The time taken to complete each subtask will be measured as well as the error rate (measured as the number of undo invocations made during task execution). Users will also be given a questionnaire to assess the systems' ease-of-use. Finally, an expert will evaluate the previz videos for correctness.

5. ETHICAL, PROFESSIONAL AND LEGAL EVALUATION

5.1 User experiments

In the user experiments, participants will be required to perform simple, non-dangerous tasks. They will be rewarded for participation, and given the choice of:

cash (R30 per subject per hour), 3D printed model (of appropriate size) or time spent using the Oculus Rift. This evaluation plan will require *ethical clearance* with the UCT Faculty of Science.

Some users report nausea when using head-mounted VR displays, including the Oculus Rift DK2. “VR may never solve its “simulation sickness” problem.” [Orland 2014] This is a potential concern for our VR participants, and must be taken into consideration when seeking ethical clearance. An appropriate course of action would be to require participants to complete an *informed consent form* and make it clear to them that they may stop the experiment at any time and still receive their reward.

5.2 Repeatability

We have an explicit goal to make our experiment repeatable. The major obstacles to repeatability are a lack of clear methodology, lack of code, lack of documentation (including build documentation) and legal impediments to redistribution [Collberg et al. 2015; Kovacevic 2007; Feitelson 2006].

This research experiment is funded in part by Triggerfish Animation Studio, however they claim no ownership of the output, including any software generated. We have approached UCT RCIPS, and they have approved our request to release the entire project under the GNU AGPLv3 license, or any later version [FSF 2007].

A copyleft license such as AGPL enshrines re-use of the code and other products of the project. This facilitates repeatability of the experiment and encourages repeatability in other work based on this codebase.

This re-use is limited in part by our proprietary dependencies, selected due to time and skill constraints. We consider these dependencies an unfortunate limitation, but will modularise the system such that non-proprietary alternatives could be developed.

To further ensure that the experiment is repeatable we must select or generate storyboards, 3D models and other assets such that they can be redistributed without legal or other impediments. We have identified *Big Buck Bunny* (2008) and related assets by the Blender Foundation as a suitable source of appropriate reusable content.

Finally, our reports may require photographs of participants using the system. These photographs will require a *model release form* (or *liability waiver*) from the subjects of the photographs, so that we can publish them. Subjects will need to be informed that under our project license their image could be re-used for any purpose.

6. RELATED WORK

6.1 Previsualisation

Lwabona [2015] has built a digital tabletop interface for collaborative previsualisation activities. This is a detailed implementation, covering a wider range of tasks to a finer degree of fidelity. By contrast, we are limiting our experiment to a core subset of tasks and a single-user interface.

Roosendaal [2015] gives an overview of a sketching tool included in Blender (which incorporates a traditional WIMP-based interface). This method requires drawing skill in addition to basic familiarity with the software application, and has not seen formal evaluation. It does however present trained digital artists with the possibility of compressing the storyboard, animatic and previz phases.

6.2 Semaphoric-tangible modalities

A tangible interface presents the user with physical objects that they can manipulate. As they manipulate the objects, they are simultaneously modifying the system's internal model [Ishii 2008]. Tangible interfaces often require at least one additional modality to account for their inherent limitations in expressive power and versatility.

Multimodal and reality-based design guidelines suggest that a companion interface should: be based in the same real-world model, be somewhat aligned with existing user activity and complement these tradeoffs [Jacob et al. 2008; Oviatt 1999; Norman 2010]. These characteristics are true for hand-based semaphores, suggesting that semaphoric-tangible interface implementations are a promising avenue for further investigation.

6.3 Model-surface registration with a commodity depth sensor

The release of the Kinect for Xbox 360 depth sensor in 2010 stimulated publications in the field of computer vision, including applications to *model-surface registration*, where a collection of known 3D points (the reference cloud) is aligned against an unknown, noisy collection of points from the depth sensor (the source cloud) [Berger et al. 2013]. Iterative closest point (ICP) is a popular algorithm well-suited to this task [Chen and Medioni 1991; Besl and McKay 1992; Salvi 2007].

The recently released K4X1 sensor offers a greater field of view, image resolution and depth resolution. [Butkiewicz 2014] When considering an implementation using a high-fidelity sensor, we note that registration based on invariant geometric features has been shown to improve the convergence rate of the ICP algorithm [Sharp 2001].

6.4 Virtual reality

There is a large body of existing research on 3D user interaction with VR, particularly with regards to 3D modelling applications. The tasks involved in 3D modelling are similar to those in previsualization, and current software limits user interaction to a 2D display combined with 2D input devices such as a keyboard and mouse. This problem has lead to 3D interaction solutions being developed that make use of VR as a more natural method of interaction. Butterworth et al. [1992] developed a 3D modelling system with a head-mounted display, supporting the choice of display arguing that by placing users in the scene they are better able to understand the 3D spatial relationships between objects. More recently, both Hughes et al. [2013] and Ponto et al. [2013] developed CAVE-based 3D modelling systems and found that users were able to create complex scenes in minutes. While there are many different ideas of how to implement 3D interfaces it is clear that a 3D user interface provides a more natural and efficient solution for 3D modelling. While this makes a 3D user interface for previsualization seem promising, there has been little research into this.

There are some tasks that pose problems for a 3D user interface. Tasks such as 2D menu navigation and numeric data entry are difficult to perform using a 3D interface [Wang and Lindeman 2014]. A new approach to mitigate these issues is to combine a 3D device with a 2D device, in Wang and Lindeman's [2014] study a 3D head-mounted VR display was used together with a tablet for 2D touch input. Another issue is 3D navigation within a scene. While a head-mounted display is good for fine view adjustments the chosen input device must allow the user to move around the scene.

7. ANTICIPATED OUTCOMES

7.1 Evaluation

Our primary output will be an evaluation of the tabletop and VR systems, highlighting their strengths and weaknesses, and their effectiveness as interfaces for generating lo-fi previz. A secondary output will be an evaluation of the applicability of the K4X1 sensor to a dynamic tabletop scenario. These evaluations will be supported by clear, repeatable methodology.

Proposed contents: overview, prototype report (already completed), tabletop literature review/ evaluation/report, registration literature review/evaluation/report, VR literature review/evaluation/report, system comparison (addendum).

7.2 Codebase

Supporting this methodology will be a reusable codebase (including all required assets) that can be used to rebuild the systems, repeat the experiments, or use as a foundation for future work, including for non-previz activities. This includes clear code architecture, repeatable build instructions and readable code.

The codebase will include:

- (7.2.1) Base classes for keyframe manipulation and timeline navigation
- (7.2.2) Tabletop previz system
  - (7.2.2.1) Tangible subsystem
  - (7.2.2.2) Registration subsystem using K4X1 sensor
  - (7.2.2.3) Input subsystem using the Myo gestural armband
  - (7.2.2.4) GUI subsystem
- (7.2.3) VR previz system using the Oculus Rift DK2 HMD

7.3 Expected results

Our hypotheses were formulated after literature reviews in related research areas, and so we anticipate that our experimental results will support them. That is, we expect that casual users will be able to use either system to produce lo-fi previz that is an accurate reflection of a predetermined storyboard, and these users will prefer the alternate interfaces to the traditional WIMP interface.

In the event that the systems do not meet this expectation, we will review on which measures they failed and review our design, implementation and experimental methodology to identify potential causes.

7.4 Success factors

We will measure success of the project on by the delivery of a ratified experimental methodology for the storyboard activity and dynamic tabletop registration, systems on which those experiments have been performed, and a downloadable resource allowing others to inspect our code and reproduce the experiment.

8. PROJECT PLAN

8.1 Key project deliverables

We have identified the following key project deliverables:

(8.1.1) Feasibility preview of each system (tabletop and virtual reality) and key subsystems (registration subsystem)
(8.1.2) Completed experimental design
(8.1.3) Completed base classes
(8.1.4) Final systems
(8.1.5) Experimental results (as individual papers, group poster and group website)

8.2 Work allocation

We have allocated work as follows:

(8.2.1) We have already performed a preliminary heuristic evaluation as a group

(8.2.2) We will program the base classes as a group (work allocation for the related subtasks still to be determined)
(8.2.3) The tabletop system will be programmed and evaluated by Katherine Rix
(8.2.4) The registration subsystem will be programmed and evaluated by David Rix
(8.2.5) The VR system will be programmed and evaluated by Joshua Ramsbottom

Each individual is responsible for the programming of their own system (or subsystem), their own experimental design and for executing their evaluations.

8.3 Resources required

8.3.1. Character assets. We require character assets that are domain-appropriate (usable in an animated movie), 3D printable, compatible with Unity and can be freely redistributed. Ideally they should also have supporting material such as storyboards. The assets for *Big Buck Bunny* meet all of these criteria.

8.3.2. Hardware. We require system-specific hardware components: a 3D printer with printing material, a depth sensor, a gestural armband and VR HMD. We also require computers capable of supporting these devices (for example, with appropriate graphics cards), and suitable for a Unity development environment.

The department has provided a MakerBot desktop 3D printer that uses the Fused Deposition Modelling method along with ABS/PLA plastic printing material. They have also supplied a K4X1 depth sensor with PC adapter, Oculus Rift DK2 headset and suitable computers. We are currently borrowing a Myo armband from a colleague, and the department will be ordering one shortly.

8.3.3. Experimental area. We will also need a location with enough space for testing and finally evaluating the systems. We have been allocated the department Makerspace (aka The Dungeon) for this purpose.

8.3.4. Code repository / development management system. We will also be using the department *GitLab* installation to manage the codebase and task allocations while the project is active. Because we will also be sharing large binary files this way we will be complementing this with a *git-annex* repository, also hosted by the department.

8.4 Risks

We acknowledge that this project has a somewhat ambitious scope and high technical risk. To this end our general approach is to make frequent, usable deliverables of incremental functionality, and anticipate downscaling individual projects if any risks cause significant delays. Specific mitigation strategies are outlined below:

8.4.1. Lack of appropriate 3D assets. In order to test and evaluate the systems, we require a consistent set of appropriate 3D assets.

Risk: Medium

Impact: High

Mitigation: We have already sourced freely licensed assets from the Blender Foundation, namely those used for their animated short *Big Buck Bunny*.

8.4.2. 3D assets are not printable. The tabletop system and registration subsystems require 3D assets that can be printed using the desktop 3D printer.

Risk: Low

Impact: Medium

Mitigation: The tabletop team have previously taken the 3D printing module and are aware of mitigation techniques including how to identify problematic models, orient objects for printing, how to add support structures and the existence of advanced balancing/remodelling techniques [Prévost et al. 2013].

8.4.3. 3D assets do not print correctly. Even if the assets are printable, 3D printing is somwhat unreliable with contemporary desktop printers. It can take time to set up the model and printer correctly and confirm that the print is successful. Despite all efforts, a particular print may fail for any number of reasons beyond the operator's control.

Risk: High

Impact: Medium

Mitigation: Starting the printing process early will allow for failures and reattempts without delaying the development process.

8.4.4. Hardware delays. Specific hardware devices are required for both systems. Some of these devices are not available locally and must be imported, thus there is a risk that there will be a delay in the delivery of one of these devices. These devices are required for development and testing, and so development may be delayed in turn.

Risk: Low

Impact: High

Mitigation: Most of the vital devices have already been obtained, and we have borrowed a Myo armband while waiting for the department device to arrive.

8.4.5. Team member permanently unavailable. One of the three team members may be unable to continue working on the project, for example due to illness or exiting the Honours program. This will impact short term deliverables (see below) and long term deliverables (whole systems / subprojects).

Risk: Low

Impact: Low – High

Mitigation: There are limited interdependencies between the two systems. If one system falls away the other system is unaffected, but the remaining team members will have a greater workload for any incomplete group deliverables. If they cannot manage this additional workload they will have to reduce the scope of their subprojects.

We can also limit the impact of such an event by ensuring that any groupwork is placed early in the project delivery schedule.

However, there is an important dependency between the registration subystem and the tabletop system. If the registration subsystem fails or is not completed in time, the tabletop system evaluation will fall back on a Wizard of Oz approach.

8.4.6. Individual delivery failure (groupwork). There are still a number of groupwork deliverables that the team depends on (including experimental design, base classes, selection and other tasks relating to the comparison WIMP system, website, poster). For a number of reasons an individual team member may fail to deliver their portion of the groupwork on time.

Risk: High (already experienced with one group member)

Impact: Medium – High

Mitigation: As mentioned above, the impact of this risk is reduced if such deliverables are placed early in the project schedule. In addition, groupwork deliverables will be unambiguously scoped and be given hard deadlines.

In the event that there is a delivery failure, this will be documented. Work will be reallocated and due credit and mark allocation adjusted accordingly.

8.4.7. Implementation takes too long. If the implementation of a system is completed late, there will not be enough time for user experiments.

Risk: Medium

Impact: High

Mitigation: In this case we will fall back to a software engineering deliverable for the late system. In other words, an evaluation of the process used to build the system, rather than of the system itself. If time permits, a more detailed expert heuristic evaluation will be performed.

8.4.8. Scope too large / too small. The scope of a subproject may end up being too large or too small. This means there may be too much work (see *Implementation takes too long* above) or insufficient work for an Honours project.

Risk: Low

Impact: High.

Mitigation: It is the responsibility of individual team members to communicate with the supervisor about the scope and progress of their subproject, and to ensure that their subproject is on the right track.

8.5 System milestones

8.5.1. Experimental design (each major system):

(8.5.1.1) Identify appropriate measures for all hypotheses
(8.5.1.2) Confirm final storyboard (based on assets and feasible transitions)
(8.5.1.3) Finalise evaluator's script
(8.5.1.4) Finalise methodology

8.5.2. Experimental design (registration subsystem):

(8.5.2.1) Identify appropriate measures for all hypotheses
(8.5.2.2) Confirm testing environment
(8.5.2.3) Finalise methodology

8.5.3. WIMP system (used for comparison):

(8.5.3.1) Identify appropriate, pre-existing WIMP system (based on functionality and known applicability to previz)
(8.5.3.2) Identify any training material required
(8.5.3.3) Confirm that, with training, the system can be used to complete the storyboard activity

8.5.4. Base classes:

(8.5.4.1) Keyframe CRUD
(8.5.4.2) Sequence store / load functions
(8.5.4.3) Timeline navigation
(8.5.4.4) Sequence interpolation

8.5.5. Tabletop system:

(8.5.5.1) Semaphoric interface
(8.5.5.2) Integration with registration subsystem
(8.5.5.3) Tangible interface (based on integration)
(8.5.5.4) GUI
(8.5.5.5) Improvements based on feedback from heuristic evaluations
(8.5.5.6) Final evaluation with 40 users

8.5.6. Registration subsystem:

(8.5.6.1) Basic registration with appropriate software libraries
(8.5.6.2) Registration of multiple objects
(8.5.6.3) Integration with tabletop system
(8.5.6.4) Final evaluation in dynamic tabletop context

8.5.7. VR system:

(8.5.7.1) More research is done in order to determine the ideal input device to be used.
(8.5.7.2) 3D navigation is implemented, making use of the headset's built-in head tracking for minor movements and the input device for larger movements.
(8.5.7.3) The first expert heuristic evaluation is performed to evaluate both the input system and the 3D navigation system.
(8.5.7.4) The 3D selection system is implemented, making use of gestures and/or the input device to allow users to select objects in the scene and menu options.

(8.5.7.5) The VR-specific floating GUI is implemented, this includes the timeline, normal menus, and floating object menus.

(8.5.7.6) The 3D object manipulation mechanisms are implemented, allowing users to perform various transformations on objects within the virtual environment by first selecting them and then using the menu system.

(8.5.7.7) The second expert heuristic evaluation is performed to evaluate the 3D selection, menu, and object manipulation systems.

(8.5.7.8) Improvements made on existing systems based on heuristic evaluations.

(8.5.7.9) Final evaluation with 40 users

8.6 Timeline

We have outlined a week-by-week timeline in *Appendix C: Delivery schedule*. This will be a living document, adjusted and kept up-to-date on the project website, with progress reports recorded every Friday morning.

REFERENCES

Kai Berger, Stephan Meister, Rahul Nair, and Daniel Kondermann. 2013. A state of the art report on Kinect sensor setups in computer vision. In *Time-of-Flight and Depth Imaging. Sensors, Algorithms, and Applications* (2013), Springer Berlin Heidelberg, 257-272.

Paul J. Besl and Neil D. McKay. 1992. A Method for Registration of 3-D Shapes. *IEEE Trans. Pattern Analysis Mach. Intell.* 14, 2 (February 1992), 239-256.

Thomas Butkiewicz. 2014. Low-cost coastal mapping using Kinect v2 time-of-flight cameras. *Oceans* – St. John's (14-19 September 2014), 1-9.

Jeff Butterworth, Andrew Davidson, Stephen Hench, and Marc. T. Olano. 1992. 3DM: A Three Dimensional Modeler Using a Head-mounted Display. In *Proceedings of the 1992 Symposium on Interactive 3D Graphics (I3D '92)*. ACM, New York, NY, USA, 135–138.

Yang Chen and Gérard Medioni. 1991. Object modeling by registration of multiple range images. In *Proceedings of the IEEE International Conference on Robotics and Automation 3*, 9-11 (April 1991), 2724-2729.

Christian Collberg, Todd Proebsting, and Alex M. Warren. 2015. Repeatability and benefaction in computer systems research. *Technical report 14-04*. University of Arizona. 68 pages. Retrieved June 12, 2015 from http://reproducibility.cs.arizona.edu/v2/RepeatabilityTR.pdf

Dror G. Feitelson. 2006. Experimental computer science: the need for a cultural change. *Manuscript*. Retrieved June 12 2015 from http://www.cs.huji.ac.il/~feit/papers/exp05.pdf

Free Software Foundation. 2007. *GNU Affero General Public License*. Retrieved June 12, 2015 from https://gnu.org/licenses/agpl.html

Sacha Goedegebure (director). 2008. *Big Buck Bunny*. Retrieved June 12, 2015 from https://www.youtube.com/watch?v=on21u4BsuK0

Cathleen E. Hughes, Lelin Zhang, Jurgen P. Schulze, Eve Edelstein, and Eduardo Macagno. 2013. CaveCAD: Architectural design in the CAVE. In *2013 IEEE Symposium on 3D User Interfaces (3DUI)*. 193–194.

Hiroshi Ishii. 2008. The tangible user interface and its evolution. *Commun. ACM* 51, 6 (June 2008), 32-36.

Robert J.K. Jacob, Audrey Girouard, Leanne M. Hirshfield, Michael S. Horn, Orit Shaer, Erin Treacy Solovey, and Jamie Zigelbaum. 2008. Reality-based interaction: a framework for post-WIMP interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 201-210.

Jelena Kovacevic. 2007. How to encourage and publish reproducible research. In *IEEE International*

Conference on Acoustics, Speech and Signal Processing (ICASSP 2007) 4, IV-1273. IEEE.

Kwegyir (Bilo) Lwabona. 2015. *AniTouch: Multitouch-based Collaborative Pre-visualisation for Computer Animation.* Master's thesis, work in progress. University of Cape Town, South Africa.

Donald A. Norman. 2010. Natural user interfaces are not natural. *interactions* 17, 3 (May 2010), 6-10.

Sharon Oviatt. 1999. Ten myths of multimodal interaction. *Commun. ACM* 42, 11 (November 1999), 74-81.

Sharon Oviatt, Rachel Coulston, and Rebecca Lunsford. 2004. When do we interact multimodally?: cognitive load and multimodal communication patterns. In *Proceedings of the 6th international conference on Multimodal interfaces (ICMI '04)*. ACM, New York, NY, USA, 129-136.

Kyle Orland. 2014. Developer cites motion sickness in delaying Oculus Rift support. *Ars Technica* (August 21 2014). Retrieved June 12 2015 from http://arstechnica.com/gaming/2014/08/developer-cites-motion-sickness-in-delaying-oculus-rift-support/

Kevin Ponto, Ross Tredinnick, Aaron Bartholomew, Carrie Roy, Dan Szafir, Daniel Greenheck, and Joe Kohlmann. 2013. SculptUp: A rapid, immersive 3D modeling environment. In *2013 IEEE Symposium on 3D User Interfaces (3DUI).* 199–200.

Romain Prévost, Emily Whiting, Sylvain Lefebvre, and Olga Sorkine-Hornung. Make it stand: balancing shapes for 3D fabrication. In *ACM Transactions on Graphics (SIGGRAPH 2013)*. 32(4), 81. ACM New York, NY, USA.

Joaquim Salvi, Carles Matabosch, David Fofi, and Josep Forest. 2007. A review of recent range image registration methods with accuracy evaluation. *Image Vision Comput.* 25, 5 (May 2007), 578-596.

Gregory C. Sharp, Sang W. Lee, and David K. Wehe. 2002. ICP registration using invariant features. *IEEE Trans. Pattern Analysis Mach. Intell.* 24, 1 (January 2002), 90-102.

Joshua G. Tanenbaum, Amanda M. Williams, Audrey Desjardins, and Karen Tanenbaum. 2013. Democratizing technology: pleasure, utility and expressiveness in DIY and maker practice. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2603-2612. ACM, 2013.

Jia Wang and Robert Lindeman. 2014. Coordinated 3D Interaction in Tablet- and HMD-based Hybrid Virtual Environments. In *Proceedings of the 2nd ACM Symposium on Spatial User Interaction (SUI '14).* ACM, New York, NY, USA, 70–79.

Received May 2015; revised June 2015; accepted June 2015

Online Appendix to:
Proposal for evaluating the effectiveness of semaphoric-tangible and virtual reality interfaces for a lo-fi previsualization activity

JOSHUA RAMSBOTTOM, UNIVERSITY OF CAPE TOWN
KATHERINE RIX, UNIVERSITY OF CAPE TOWN
DAVID RIX, UNIVERSITY OF CAPE TOWN

A. REVISION NOTES

Revision 1. First submission.
Revision 2. Major revision of all sections. Added “how does it work?” comic.
Revision 3. Added revision notes, delivery schedule. Other minor improvements.
~~Revision 4 (planned). Add AGPL license notice. Use hi-res comic images.~~

B. TABLETOP PREVIZ: HOW DOES IT WORK?

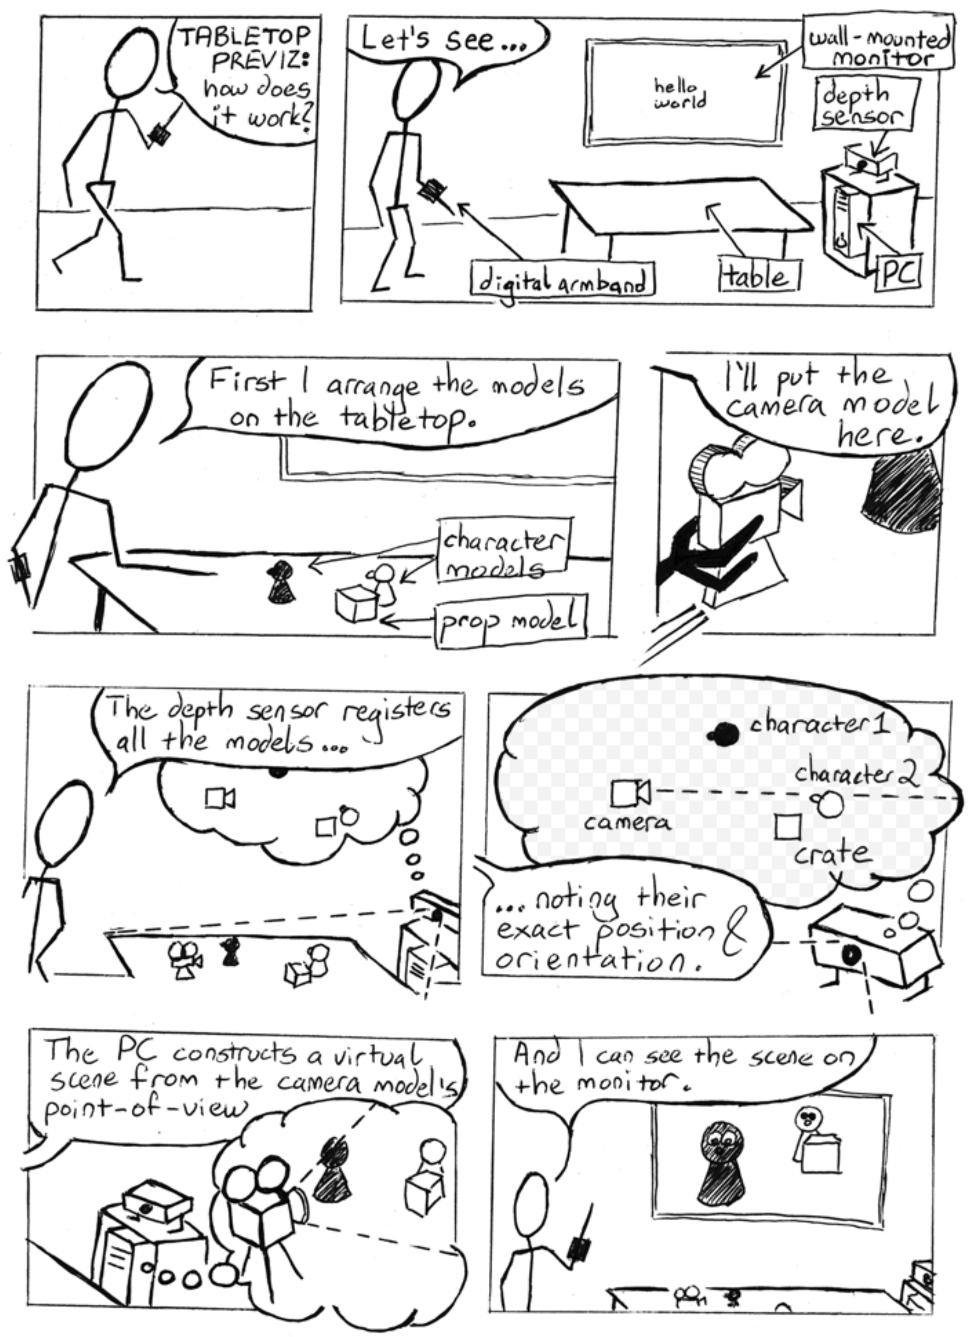

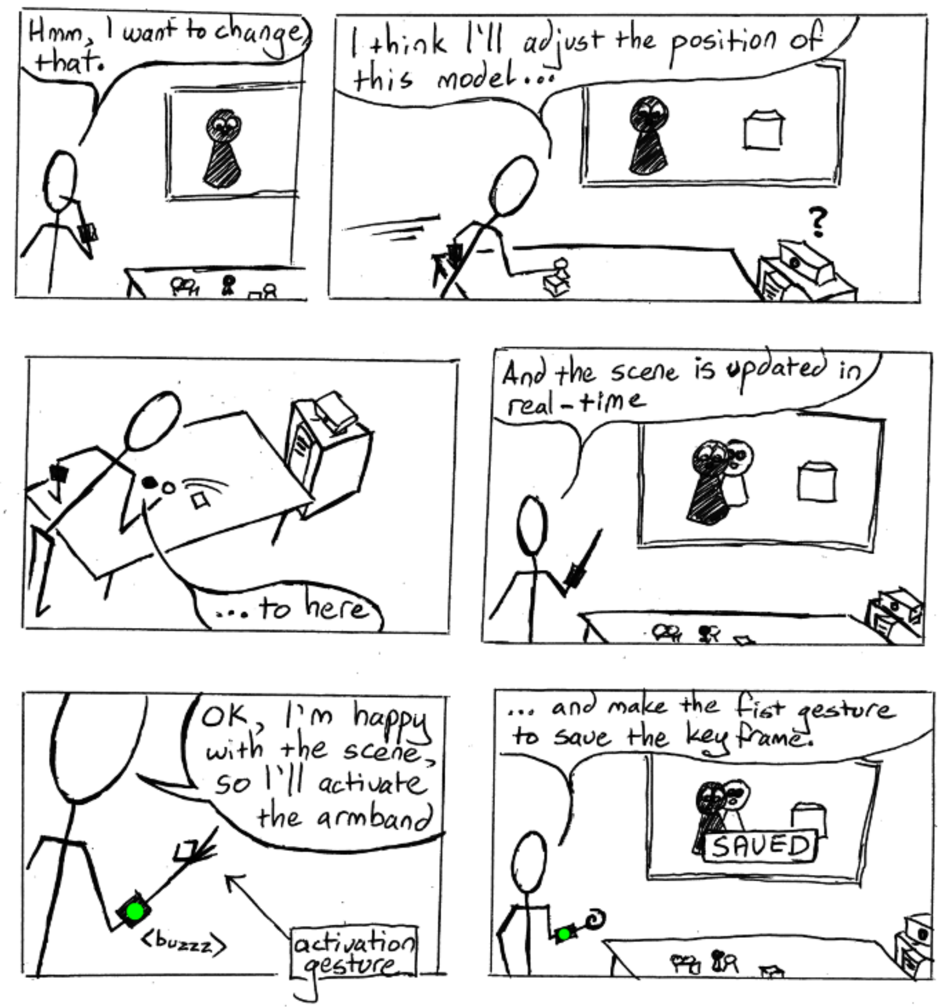

Tabletop previz: how does it work?

C. DELIVERY SCHEDULE

June 2015

Week	Tabletop	Registration	VR	Other
2015-06-01	Revise proposal and do website (all)			Block 2 Exams
2015-06-08 *	-	-	-	Block 2 Exams
2015-06-15	- Myo / Unity integration	- Kinect / Unity integration	- Oculus / Unity integration	-
2015-06-22	- Experimental design - 3D printing	- OpenCV integrat. - Git + git-annex - 3D printing	- Decide input device - Device integration	Triggerfish
2015-06-29	- Interaction design - Task analysis - 3D printing	- Rudimentary ICP - 3D printing	- Scene navigation with Oculus	Triggerfish

Handins (marked with *):

- Revised proposal (June 15)
- First website (June 15)

Other commitments:

- Exams for Block 2
- Triggerfish internship (and informal ethnographic study)

July 2015

Week	Tabletop	Registration	VR	Other
2015-07-06	Base classes (collaboration)			Triggerfish
2015-07-13	- Design review (for individual demo)	- System review (for individual demo)	- System review (for individual demo)	Triggerfish
2015-07-20 *	- Myo control of timeline for 3D scene	- ICP implementation	- 3D object selection	Block 3
2015-07-27	- Heuristic evaluation - Fixes	- ICP implementation	- 3D object selection	Block 3

Handins (marked with *):

- Individual feasibility demos (July 20)

Proposed handins (not confirmed):

- Background/theory submission (July 24)

Other commitments:

- Block 3 modules

August 2015

Week	Tabletop	Registration	VR	Other
2015-08-03	- Preliminary sensor integration	- First experiment (single model)	- 3D object selection	Block 3
2015-08-10	- Sensor integration	- Experimental design - Sensor integration	- Floating menus	Block 3
2015-08-17	-	-	- 3D object manipulation	Block 3 Exams
2015-08-24	-	-	- 3D object manipulation	Block 3 Exams
2015-08-31	- Heuristic evaluation - Fixes	- Build instructions - Repeatability	- Expert heuristic evaluation	-

Handins (marked with *):

- None

Other commitments:

- Block 3 modules
- Exams for Block 3

September 2015

Week	Tabletop	Registration	VR	Other
2015-09-07	- System review - Build instructions - Repeatability	- Experiments	- User evaluation	-
2015-09-14	- User evaluation	- Experiments	- User evaluation	-
2015-09-21	- User evaluation	- Stats & diagrams	- Writeup (outline)	-
2015-09-28 *	- Writeup (outline)	- Writeup (outline)	- Writeup (outline)	-

Handins (marked with *):

- Complete outline (2 October 2015)

October 2015

Week	Tabletop	Registration	VR	Other
2015-10-05	- Stats	- Writeup (draft)	- Writeup (draft)	-
2015-10-12 *	- Writeup (draft)	- Writeup (draft)	- Writeup (draft)	-
2015-10-19	- Writeup (polish)	- Writeup (polish)	- Writeup (polish)	-
2015-10-26 *	-	-	-	-

Handins (marked with *):

- Individual report, draft (16 October 2015)
- Individual report, final (26 October 2015)