

WeCurate: Designing for synchronised browsing and social negotiation*

Leila AMGOUD, Roberto CONFALONIERI, Dave DE JONGE, Mark d'INVERNO, Katina HAZELDEN, Nardine OSMAN, Henri PRADE, Carles SIERRA, and Matthew YEE-KING

¹ Goldsmiths, University of London

² Institut de Recherche en Informatique de Toulouse

³ Artificial Intelligence Research Institute (IIIA), Catalonia

Abstract. WeCurate is a shared image browser for collaboratively curating a virtual exhibition from a cultural image archive. This paper is concerned with the evaluation and iteration of a prototype UI (User Interface) design to enable this community image browsing. In WeCurate, several remote users work together with autonomic agents to browse the archive and to select, through negotiation and voting, a set of images which are of the greatest interest to the group. The UI allows users to synchronize viewing media, assists navigating the cultural institutions extensive database and users negotiations about which images should be added to the group's image collection. The objective of WeCurate is to provide a platform for a sociocultural experience, combining the actions of autonomic agents and users to facilitate decision-making, and Electronic Institutions to articulate and structure the behavioural norms of groups.

1 Introduction:

WeCurate is an online image browser which synchronises the images presented to multiple users, so they are seeing the same image at the same time. The purpose of the system is to provide an environment for mediating socio-cultural experiences via engagement with a Cultural Institution's (CI) artefacts. The common goal for this group activity is to compile a subset of images from the CI image database to create a collaborative virtual exhibition. This activity mirrors the task of museum practitioners, and enables the creation of a collection that is personal to the collaborative effort. WeCurate makes use of autonomic agents and Electronic Institutions (EI) to aid the agreement about which images are discussed and saved to the collection. This paper presents the findings from the evaluation of the first prototype, WeCurate V1, and outlines the plans for the next iteration of design.

The system intends to both connect new audiences to the CI's artefacts, and enable existing audiences to engage with CIs databases in a novel way. This activity may provide access for those unable to physically visit the CI, provide an

* AT2012, 15-16 October 2012, Dubrovnik, Croatia. Copyright held by the author(s).

awareness of the CIs artefacts prior to visiting, or re-enforce the experience after visiting. Additionally, WeCurate aims to facilitate access to the large number of artefacts in the CIs archives, as many items do not feature in current exhibitions. Prominent CIs in the UK have online databases that contain thousands of images; the British Museum for example, has 634,418 objects in their archive [20] and is constantly expanding. Recursively, the design, development and evaluation of WeCurate also functions as a vehicle to better understand the role of agents to assist social negotiation and data retrieval.

Physically visiting a CI is an experience augmented by others, it is an activity shared with families, friends and education groups [12, 17] Sociality enriches the experience, for example, more informed members of the social group frame and scaffold the knowledge associated with the CI exhibits [7], and peer groups will attempt to make sense of information together, drawing on personal histories [1]. Consequently, the design of WeCurate needs to be mindful of how community members will engage with the environment, the task, and the dynamics of their collaboration. The design approach pays careful attention to the social interaction and representations of social action, to generate a strong sense of social presence among the users. Social presence relates to the awareness of others, and to some extent the autonomic agents' actions, intentions and judgements in the mediated space. This is coupled with an understanding that others share a mutual attention, mutual comprehension and interdependent behaviour [3].

The version of the WeCurate system which is trialled in this paper, is an initial prototype which aims to provide a simple, synchronous shared image browsing experience for two users. The system is implemented for iOS running on iPad hardware; this platform was chosen as it is a highly portable, rich media capable device with a multitouch screen and powerful processors. Also, iOS has built in libraries which facilitate peer to peer networking and rapid GUI construction. The software presents the user with a simple interface showing an image along with its associated tags. There are *next image* and *save image* buttons, and a text field through which the system communicates events to the users such as *The other user hit the next image button*. The user can express a like / dislike preference for the tags, and can also see the current tag preferences of the other user in a duplicate tag list next to theirs. Once both users tap the next image button, the system selects the next image for them by merging their tag preferences and using the result to search the image database for the image whose tags best fit their preferences. As such, it might be classified as a simple online image recommender, though image recommendation is not the main focus of this work; rather socially browsing images is, in which image recommendation plays a supporting role.

The system is modelled using several agents with distinct roles operating in a simple, peer to peer electronic institution (EI). A *media archive agent* processes requests for images. A *deliberating agent* collects both users' tag preferences during the session such that it can merge these preferences and select the best match for them from the database when the users request the next image. A *user assisting agent* represents the user in the EI, receiving user interface events

and passing them on to other agents as appropriate, for example, it passes user tag preferences to the deliberating agent. In order to mediate the messaging between these agents, a lightweight EI has been implemented. It controls which messages can be sent by which agents and when. Modelling the software in this way makes it more complex than it needs to be to solve the two user, image browsing problem; typically, electronic institutions are more appropriate for modelling complex interactions between many agents, playing a variety of roles in a variety of scenes. However, the motivation for modelling the system in this way is simple: the two user image browser is the first prototype and future systems will enable much richer user experiences, including more users and more numerous and complex agents and the agent/ EI model is designed to scale to these more complex tasks. The EI has been implemented in a peer to peer fashion where components of the EI are executed on both devices rather than on a centralised server. The motivation for implementing a peer to peer EI rather than a server based EI was an interest in the development of new, potentially more robust types of EI infrastructure.

2 Related work

This section highlights principles which frame the design of socio-cultural experiences via social presence. Enabling cultural experiences to be social offers potentially significant gains to understanding the meaningfulness of cultural artefacts. Working with others can extend the competencies of individual group members, as the gestalt of a collective ability is capable of exceeding that of individuals working in isolation [2, 14, 15]. This is not a total homogeneity, individuality remains intact, and personal traits will affect the shape of group characteristics. Each individual adds to the behavioural diversity of the group by continually adjusting his position in the network of interactions that forms the group according to its dynamics [15].

To enable social collaboration, it is paramount to the design of WeCurate that there is a sense of social presence among the group. Social presence in online social groups can emerge around a shared social identity and joint purposefulness, rather than physicality [16]. To coordinate group actions, the interface must communicate the user and agents actions in relation to others, to provide a context for individual activity [6]. The synchronisation of the activity in WeCurate should create a sense of collective experience, and therefore should contribute to social presence. An understanding that collaborators have the same view of the activity should make users accountable for their own actions [8]. Further, there is an expectation that arguing will be stronger, more impulsive, with more counter argument in a synchronised system, as users have less time to deliberate their response [19].

The current design of WeCurate has selected tagging as a means of identifying and communicating the semantic information about the images. Collaborative tagging, classification according to specific social groups appears to be a popular activity as a means of expressing opinion and categorising community content.

Collaborative tagging is pivotal to community websites such as Flickr, del.icio.us, and 43 Things. Tags tend to succinctly denote media content in terms of who / what it is about, what it is, qualities and characteristics (i.e. funny, stupid) according to the tagger's opinion [11]. As such, tags are not only associated with the items, but also with the users who created them [13, 18]. Though not a primary objective, there is some motivation from the CI's involved in this project to crowdsource data about their artefacts, to produce domain specific folksonomies that will assist in gaining a perspective on visitors' understanding and appreciation of their exhibits. As the system develops and is used in the real world, user generated tags will also help search for more appropriate content, as even subjective tags have wider usefulness. What is funny to one user is likely to be funny to another [11]. There are some anticipated issues with inferring meaning from tags, as language can be indistinct, for example, the use of homonyms and possible synonyms to describe the same objects [11].

Creating a sense of community and connectedness, even if it is ephemeral, is essential for a socio-cultural experience to emerge in the context of our intended system. The design framework for WeCurate is therefore constrained by the need to communicate social action and enable community members to identify and define a CI's artefacts in their own terms.

3 Study

The evaluation of WeCurate V1 utilised qualitative techniques, relying on researcher observations and the participants self-reporting behaviours and opinions, to reveal potential usability and conceptual issues with the task, and social cohesion in the collaborative activity. Observations of the task paid close attention to the communication and the dynamic between co-located participants to uncover emerging social practices. Further, attention was paid to coordinating behaviours, as awareness and responsiveness to others activity within the mediated environment may be both necessary and indicative of the movement to higher levels of co-presence [4]. As previously stated, WeCurate V1 provided a simple UI, focusing on the core tasks. When using the system, users were asked to express a *like* preference of overall image, and state which tags were *liked* or *disliked* in association with this image. The image view was synchronised, and partners tag preferences were displayed.

3.1 Method

This was a small trial with eight participants, all adults tested in co-located pairs. This is clearly not a large enough sample to denote definitive social trends, but served as an indicator of how to move forward from the first iteration of design. With one exception, the participants were all students at Goldsmiths College. Friends were sourced, rather than random pairings, as it was assumed this group were the most likely to share images on and off line.

Each trial lasted approximately 45 minutes, and consisted of three activities: (i) paper-based card sorting, (ii) use of WeCurate V1, and (iii) interviews. In the short introductory task, the participants were simply asked to collaboratively sort the images into *like* and *dislike* piles. This was a purely speculative exercise, which was designed to be a low tech version of the technologically mediated activity. This task was targeted at investigating discussions, sharing and the organisation of visual images. Also observing behaviour here could help rationalise actions during the main task. In the main task, participants were requested to use WeCurate V1 with their partner for 20 minutes. Finally, three key areas were discussed in loosely structured interviews: (i) a brief explanation and comparison of current practices of sharing image with friends online, (ii) the use, conceptual understanding and opinion of our system and (iii) ideas for expansion of features, functionality and visualisation of the system's and users' actions.

3.2 Analysis

The analysis cross referenced three sources of data from each activity:

- Observations during the trial, ad hoc during task and from video recordings for closer examination
- Interviews
- Log of tag selection

The observations and interview data from the trials considered the users' opinions and engagement in the collaborative task (viewing and tagging), and their awareness of the agents selecting media according to explicit preferences. Analysis looked for indicators of social presence in the participants' reference of the media (i.e. gesturally, verbally), the efficiency of their language which would signify a high degree common ground, and therefore a shared conceptual space [5].

Aside from the simplified interface of WeCurate V1, there were two significant differences between the prototype version and the plans for the subsequently versions which needed to be accounted for in the analysis. Firstly the image collection, the trial uses a selection of images taken, with their metadata, from Flickr. The image collection planned for the eventual use is to be sourced from a CI in the London area. The second difference is the context of use. In the trials the participants are situated in the same location, although it is not yet finalised, it is likely that the WeCurate system will be used by remote or mobile users.

4 Results

In this section we present the findings from the user trials. The data sources the tag selection in the log, observations and interview data - were aggregated and categorised into interaction features and social attributes, rather than considered

in isolation. There were some usability issues resulting from the simple UI of WeCurate V1, which have not been reported as they are outside of the scope of this paper. The results outlined here relate to three pivotal aspects of the intended design: the role of the tags in defining and describing the image, the users' awareness of the agent's actions, and a description of the communication and social behaviours that emerged in testing.

4.1 Tags

All of the participants reported that the tags were insufficient, and therefore at times they struggled to express a positive or negative preference via the interface. As anticipated, there were some misunderstandings over the meaning or relevance of the tags which further contributed to difficulties in expressing a preference. Each image had six to nine tags attached to the image, which in this situation appeared to be too limiting. However, as is shown in Extract 1 below, one participant commented that was not just the quantity, but the quality of the tags that hindered the experience.

Extract 1. Interview.

I think that most of them, there are not enough. Most of them are abstract pictures, and I like the idea, I like the composition or the angle of the picture. Not just because of the colour, or the situation, like ice, frozen.

Important descriptive elements of the image which were attractive for the viewer were not present in the current tag list. There was also some repetition of meaning, where the tags were not distinct enough (i.e. bike and bicycle, or fog and mist), which made expressing a preference a somewhat arbitrary exercise.

4.2 Awareness of the algorithm

When the participants were given instructions to use the system, the filtering agent was not fully explained. Despite this lack of information there was a realisation that the images presented in WeCurate V1 were determined by the tag selection. Though awareness was apparent, its presence was not overwhelming, and it was not entirely well received. As demonstrated in the extract below, the similarity between the images could become tedious.

Extract 2. Interview.

I disliked the pictures more. I become tired of the pictures because some of them are similar to some extent and I became tired of this.

The participant from the Extract 2, noticing that similar images were being selected, later reported that she purposefully modified the tags she was selecting in an attempt to vary the images.

4.3 Social Interaction

There were multiple channels of communication observed during the evaluation. Primarily the participants relied on direct communication, expressing their ideas verbally, which was typically supported by non-verbal gestures. The participants also communicated via the interface, by expressing a preference with the tag selection and the *Like Image* option, which was either visually or textually represented on their partner's screen.

There was a radical difference when comparing the percentage of talk time in the paper-based task to using the WeCurate V1 application, despite no change in the physical setting as one task flowed on from the other. On average, 93% (SD=4.4) of the time in the pre-task activity was spent in conversation. This was to some degree expected, as each pair was required to reach a consensus about which images to keep and which to discard. For all the groups, there were gaps of no more than 2-3 seconds during the paper-based activity. The reduced talk time during the WeCurate V1 activity was substantial, with more variation, averaging 36% (SD=21.4) of the activity. Aside from being a much longer task where conversation may have naturally tailed off, WeCurate V1 did not necessitate an agreement; rather it made each party aware of the decisions made by their partner. Consequently, there was no forced social interaction when using the application, it was instead determined by the interest of each party in each other's expressions of preference.

4.4 Social presence

There appeared to be only slight awareness of partner's action when reporting the experience of using the WeCurate V1. One group noted that they were only 'vaguely' aware at the beginning, others reported much the same: 'I guess at the beginning perhaps more, but then you just get used to it and get in the flow'. So despite the persistent presence of the dynamic information of their partner's activity, this awareness appeared to diminish over time.

When using the WeCurate V1 application, the participants' language was very efficient with no observable breakdowns in communication, reflecting the effectiveness of the synchronised screens in a co-located setting. In Extract 3, the efficient communication between the participants indicates that, in spite of the reduced talk time for this group, there was an awareness of their partners' activity via the peer-to-peer system. Here, the participants have a shared reaction to content in WeCurate V1, without communicating what they are referencing.

Extract 3. Transcript

Sp2. Oh-

Sp1. ((Laughs))

(13.)

Sp2. Hmm. Ah.

Sp1. That (disturbs) me.

Sp2. ((Laughs)). It's nice.

The succinct language observed in this verbal exchange was possible as the participants assumed they shared the same view. Consequently, there was no need for the speakers in Extract 3 to identify the referent, or visibly look for affirmation - knowing a friend's potential attitude, coupled with the synchronised viewer in the WeCurate V1 interface, enabled a common conceptual understanding between the participants with no effort.

4.5 Tagging as a social action

The examination of the trial data looked for additional evidence of exchanges via WeCurate V1 in the participants' tag selection. More specifically, whether there was a correlation between the selection of tags i.e. if the selection of a tag by one participant prompted their partner to select to the same tag. When all four sessions were aggregated, the log data revealed that 17% of tag selections were linked, which does not show a significant trend. Whether this is because their attention to their partner's tag selection diminished through the course of the evaluation, or that one user's tagging did not provoke a reciprocal action from their partner, is indistinct.

However, although there was not overwhelming evidence, an awareness of user's tagging was reported as influencing some aspects their partner's behaviour. During one interview, a participant mentioned that seeing her partner's tag preference selection on her screen made her re-examine the image for a detail she may have overlooked. Additionally, as Extract 3 shows, back channel responses were triggered by understanding their partner's action / reaction via the WeCurate V1 interface. Showing that a degree of communication was occurring non-verbally via the peer-to-peer interface. When they did take notice, the participants reported that they found seeing their partners' actions enjoyable. Extract 4 below shows, there were also indicators that tagging was a means to better understand your partner.

Extract 4. Interview.

Maybe you get interested in the other person? Or the other way round, if you always tag differently. I think you would, like I said I didn't always compare our taggings but I guess if you don't know the other person you might become more aware of the comparison.

The participant believed she was able to make inferences about their partners character and intentions, even speculatively in circumstances where their partner was a stranger. Finally, although a minor point, one participant commented at the beginning of the WeCurate V1 task:

Extract 5 Transcript.

Sp2. It's a bit creepy. Seeing it. (((Laughs)))

Sp1. (((Laughs)))

This comment was in reference to seeing their partner's tagging actions on their own screen. This is not perhaps the most positive expression of social

presence, as the participant was slightly unnerved, but observing the actions of his partner via his interface did prompt the impression of someone else occupying the same space.

5 Discussion of WeCurate V1 results

Issues relating to the communication between the participants were prevalent in the observations and the feedback. The impact to talk time is a cause for concern, as the mediated task did appear to obstruct conversation in a co-located setting. However, there was clearly some communication via the system which arguably reduced the need to talk, and which suggests the WeCurate V1 interface provided a degree of shared social presence and identity.

The metadata tags were a key part of the social experience engineered by WeCurate V1. They were a means by which peers could, via the interface, highlight pertinent aspects of the image and communicate their opinions. Further, by viewing a peer's selection, they could draw attention to content within the image. The meaningfulness and relevance of the tags in this trial caused some confusion, and at times limited the deconstruction of the visual material. Not providing the right tags could potentially undermine the tasks in WeCurate. This is not something we will necessarily have complete control over if sourcing large collections from CI, but if we can get these right, we are likely to be able to provoke more social engagement.

There were also issues with the algorithm presenting similar images. If the filtering is not carefully handled, we run the risk of reducing opportunities for spontaneity and the range of material presented to the user. There is unquestionably mileage in filtering images to match user's tastes, particularly with weighty databases, but there is also evidence here that viewing a collection of similar images could be tiresome, as the users expect more variation and they have no autonomy to direct their browsing.

The feedback from the trial clearly showed that visual representations of their partner's action needed to be more apparent as the participants reported their attention reduced during the course of the activity. Getting this feature right is especially crucial if WeCurate is used in a remote context, as there was some communication occurring via broadcasting the user's actions from within the interface.

6 Ongoing and future work

The next step for the WeCurate application is to implement the changes which emerged in the evaluation, and to continue a process of iterative testing as the system develops and shifts to new contexts of use. In the short term, this involves expanding the peer-to-peer network to multi-user, and the connection to an image database provided by a real CI. As discussed, there are some ongoing concerns which will need to be resolved in development and further trials, namely viewing options, communicating action and negotiation, the similarity problem,

and controlling the quality of the metadata. For this last point, some solutions may emerge as we begin to engage with the CI, though it unlikely to be something that can truly be resolved without some human intervention.

6.1 Designs for WeCurateV2

WeCurate V2 will implement some changes which are thought to increase subjective and social engagement. Firstly, the activity itself will be re-structured into three phases: suggestion (is this image interesting enough to discuss?), discussion (should this image go into the collection?) and voting (definitive yes / no). It is the discussion scene that will allow the group to explore the conflicts and consensus, prior to electing the image be saved. The image and tag preference will have a scale to express a degree of preference, rather than a binary like/dislike option. In the absence of controlling metadata vocabularies, and to enable a more specific and personal discussion of the images, users will be able to post their own tags, which can then be voted on by the group. Additionally, a text comment feature will be included for opinion communication not possible with short tags. In terms of selecting images for the group to discuss, the algorithm will continue to select images from the database, but agents will be tasked with selecting images that are of interest to the group based on a more complex aggregate of users action. Actions such as expressing a preference, zooming the images, manipulation and the creation of tags, will be inferred by agents as interest among the group. Also, negative as well as positive responses will be considered, which should diversify the selection, and acknowledges that an engaging discussion can be had around a negative opinion. The role of the agents to mediate the synchronisation of group activity will also be extended to determine whether there is enough interest in the current image to move the group to a discussion or to vote.

The interface needs to enhance representations of social and agent actions in order for users to contextualise their actions with others. WeCurate V2 plans to implement the principles of social proxies, which suggests a shared visualization that shows the presence and activities of participants [9]. This includes a third person view of self-action so the user can see his activity in relation to others, and therefore can begin to make inferences about the activities of others [10]. In the next version, all the actions will be public (i.e. comparing images, zooming, add image to library, chat etc). Seeing the interactions of others would enable users to realise there was activity happening even if they were not part of the discussion, and may prompt interventions. This option would arguably promote a stronger sense of community even from those who are not actively participating. There is also a problem with users failing to consistently attend to this information throughout the activity, so it may be advantageous to attach auditory or graphical triggers to draw the user's attention to their partner's action.

The development of WeCurate V2 will be subject to further field trials to assess whether the system is fit for purpose, and to continue to evaluate emergent social behaviours. To keep discussions vibrant, yet manageable for the user, we envisage small groups (8 -12 users) using this system, with multiple groups

working simultaneously. With use, it is hoped that the selection based on meta-data will be increasingly more accurate, as this data can be stored and then re-used by the agent to locate images in the selection scene. The evaluation of WeCurate in situ will collect data via existing methods (observation and self reporting) and extend the automatic logging to include more contextual information about users actions (preferences, tags, time, social interactions) and the actions of autonomic agents.

7 Conclusion

It is hoped that WeCurate will be able to provide an environment where an online community can explore, discuss, and argue in response to identifying and deconstructing cultural artefacts. We aim to achieve this by enabling an online space where users can synchronise viewing CI images, coupled with a collaborative activity which asks them to classify and assess the artefacts. This paper has outlined the details of this work in progress, and has shown how certain aspects of synchronisation, collaborative tagging and agent mediated decision making can contribute towards the social presence required for this experience to be considered successful. Further, it has outlined which features and functionality need to be modified or fine-tuned for the design to progress.

References

1. Ash, D. Negotiations of Thematic Conversations About Biology. In G. Leinhardt and K. Crowley & K. Knutson (Eds.), *Learning Conversations in Museums* (pp. 357-400). Mahwah: Lawrence Erlbaum Associates. (2002).
2. Barsalou, L.W. Grounded cognition. *Annual Review of Psychology*, 59, pp. 617- 645. (2008).
3. Biocca, F. and Nowak, K. Plugging your body into the telecommunication system: Mediated Embodiment, Media Interfaces, and Social Virtual Environments. In D. Atkin & C. Lin (Eds.), *Communication Technology and Society*. (2001).
4. Biocca, F. and Harms, C. Defining and measuring social presence: Contribution to the Networked Minds Theory and Measure. n F.R. Gouveia, and F. Biocca (Eds). *Proceedings of the 5th International Workshop on Presence*. (2002).
5. Clark, H.H. and Brennan, S. A. Grounding in communication. In L.B. Resnick, J.M. Levine, & S.D. Teasley (Eds.). *Perspectives on socially shared cognition* . Washington: APA Books. (1991).
6. Dourish, P. and Bly, S. Portholes: Supporting awareness in a distributed work group, *Proceedings of ACM CHI 1992*, 541-547. (1992).
7. Ellenbogen, K., Luke, J. and Dierking, L. Family Learning Research in Museums: An Emerging Disciplinary Matrix? *Science Education*, 88(Supplement 1), S48-S58. (2004).
8. Erickson, T. Designing Visualizations of Social Activity: Six Claims. *The Proceedings of CHI 2003: Extended Abstracts*. In press. New York: ACM Press, 2003.
9. Erickson, T. A Social Proxy for Collective Search, *CSCW 2010 Workshop on Collaborative Information Seeking*. (2010).

10. Erickson, T., Halverson, C., Kellogg, W.A., Laff, M. and Wolf, T. "Social Translucence: Designing Social Infrastructures that Make Collective Activity Visible." CACM (Special issue on Community, ed. J. Preece), Vol. 45, No. 4, pp. 40-44. (2002).
11. Golder, S. and Huberman, B.A. Usage Patterns of Collaborative Tagging Systems. *Journal of Information Science*, 32(2). 198-208. (2006).
12. Hindmarsh, J., Heath, C., vom Lehn, D. and Cleverly, J. Creating Assemblies in Public Environments: Social Interaction, Interactive Exhibits and CSCW. *CSCW* 14(1): 1-41 (2005).
13. Krause, M. and Aras, H. Playful tagging: folksonomy generation using online games, 1207-1208. In *WWW '09 Proceedings of the 18th international conference on World wide web*. (2009).
14. Masciotra, D., Roth, W.M. and Morel, D., *Enaction*. Sense Publishers. (2006).
15. Maturana, H.R. and Francisco V. *The tree of knowledge: The biological roots of human understanding*. Boston: Shambhala. (1992).
16. Rogers, P. and Lea, M. Social presence in distributed group environments: The role of social identity. *Behaviour & Information Technology*, 24, 151-158. (2005).
17. Sanford, C., Knutson, K. and Crowley, K. "We Always Spend Time Together on Sundays": How Grandparents and Their Grandchildren Think About and Use Informal Learning Spaces. *Visitor Studies* 10:2, 136-151. (2007)
18. Tso-Sutter, K.H.L., Marinho L.B. and Schmidt-Thieme, L. Tag-aware Recommender Systems by Fusion of Collaborative Filtering Algorithms. (2008)
19. Veerman, A.L., Andriessen, J.E.B. and Kanselaar, G. *Collaborative Learning through Computer-Mediated Argumentation*. (1999)
20. British Museum Collection: http://www.britishmuseum.org/research/search_the_collection_database.aspx