# D3.1

# Identification of technologies and useful software

2016-10-28

Antonios Liapis, Martín López Nores, Alberto Gil Solla, Yannick Naudet, Ioanna Lykourentzou, Louis Deladiennee, Sylvain Castagnos, Amedeo Napoli, Aggeliki Antoniou, Nikolaos Platis, Joe Padfield, Kalliopi Kontiza, Antonios Bikakis, Andreas Vlachidis

This deliverable, grounded on a review of the state-of-the-art in the pertinent areas, aims to gather knowledge and references as to what technologies, algorithms and software may be reused, adapted or taken as a reference in the implementation of the modules offered by the CrossCult platform to support the creation and execution of applications for all the targeted users: museum curators and experts, administrators, data scientists, experience designers and museum/city visitors. The document takes input from the design of the pilots described in D2.1 (*"Pilot specifications"*).

www.CrossCult.eu

| | |
|---|---|
| **Project acronym** | **CrossCult** |
| Full title | CrossCult: Empowering reuse of digital cultural heritage in context-aware crosscuts of European history |
| Grant agreement number | 693150 |
| Funding scheme | Innovation Action (IA) |
| Work programme topic | H2020-REFLECTIVE-2014-2015/REFLECTIVE-6-2015: Innovation ecosystems of digital cultural assets |
| Project start date | 2016-03-01 |
| Project duration | 36 months |
| Workpackage 3 | Technological modules |
| Deliverable lead organisation | University of Malta |
| Authors | Antonios Liapis (UOM) Martín López Nores (UVIGO) Alberto Gil Solla (UVIGO) Yannick Naudet (LIST) Ioanna Lykourentzou (LIST) Louis Deladiennee (LIST) Sylvain Castagnos (LORIA) Amedeo Napoli (LORIA) Aggeliki Antoniou (UOP) Nikolaos Platis (UOP) Joe Padfield (NG) Kalliopi Kontiza (NG) Antonios Bikakis (UCL) Andreas Vlachidis (UCL) |
| Reviewers | Marta Redondo (GVAM) Jaime Solano (GVAM) |
| Version | 1.0 |
| Status | FINAL |
| Dissemination level | PU: Public |
| Due date | M8 (31/10/2016) |
| Delivery date | 28/10/2016 |

# Table of Contents

# List of Tables

# 1. Introduction

## 1.1.  CrossCult overview

CrossCult (www.CrossCult.eu) is a three-year H2020 research project, started in March 2016 and developed by 11 European institutions and 14 associated partners from the areas of computer science, history and cultural heritage. The goal of the project is to spur a change in the way European citizens appraise History, fostering the re-interpretation of what they may have learnt in the light of cross-border interconnections among pieces of cultural heritage, other citizens' viewpoints and physical venues. To this aim, the project will integrate innovations in three axes:

- First, maturing technologies of <u>knowledge modelling</u> and <u>semantic reasoning</u> will be harnessed to create **an extensible knowledge base** that makes the connections explicit across an unrestricted set of repositories of digital cultural heritage resources. The interconnections may derive from common facts (e.g. two assets belonging to the same or overlapping eras or places) and from crosscutting, transversal concepts that serve to view societies rather than isolated characters or events (e.g. social classes and gender issues, folklore and rituals, war or the struggle for water, food or mineral resources).

- Second, **an open technological platform** will be developed to facilitate the creation of interactive applications for individuals and/or groups to appraise the connections with the aid of recent innovations in the areas of <u>virtual and augmented reality</u>, <u>micro-augmentations</u>, <u>geolocalisation</u> and <u>sporadic social networks</u>. Besides, the platform will provide supporting tools to aid in the work of the other users behind the scenes, namely museum curators and experts, data scientists, experience designers and administrators.

- Finally, state-of-the art techniques of <u>content adaptation</u> and <u>personalisation</u> will be used to create narratives for the interactive experiences that maximise **situational curiosity and serendipitous learning**, taking into account the cognitive/emotional profiles of the participants as well as temporal, spatial and miscellaneous features of context (e.g. calendar events, weather, congestion points in a venue, or popular news in the participants' social media).

The objectives of the project can be summarised as follows:

- **[Humanities research objective]** *Meta-history research from pilot experiences with narratives built from cross-border connections and crosscutting topics.* To develop pilot experiences that investigate the potential of situational curiosity and serendipity to increase the retention of historical facts linked by cross-border connections or crosscutting topics, gaining insight into the question of how the same facts may be interpreted differently from different social realities and by individuals with different cognitive/emotional profiles (meta-history).

- **[Innovation objective]** *Multi-level, cross-repository interconnection of venues and digital cultural heritage resources.* To create a semantic knowledge base that interrelates an unrestricted set of (existing and future) digital cultural heritage resources and venues across different repositories, on the grounds of common properties or crosscutting, transversal concepts.

- **[Innovation objective]** *New technologies for smart venues and cities.* To assess the impact of state-of-the-art technologies of <u>geolocalisation</u>, <u>micro-augmentations of reality</u>, <u>social networking</u>, <u>content adaptation</u> and <u>personalisation</u> in mobile edutainment apps for smart cities and smart venues.

- **[Innovation objective]** *Personalised and context-aware experiences.* To automate the generation of narratives and the composition of digital cultural heritage resources in order to deliver meaningful interactive experiences to individuals and groups, taking into account their cognitive/emotional profiles, as well as temporal, spatial and miscellaneous features of context.

- **[Exploitation objective]** *Business models for new cultural experiences.* To design business models and plans for the exploitation of the project results, assessing the viability and sustainability of the proposed knowledge base, technological platform and interactive experiences in collaboration with a new network of researchers, scholars, ICT professionals and specialists of digital heritage.

## 1.2.  Context of the deliverable

This Deliverable aims at gathering knowledge and references about what technologies, algorithms and software may be reused, adapted or taken as a reference in the implementation works of WP3. The Work Package is about developing the modules that will be offered by the CrossCult platform to support the creation and execution of applications for all the targeted users. Therefore, the Deliverable relates mainly to the aforementioned "*innovation objectives*" and to the technologies that appear underlined in Section 1.1.

The classification of technologies has been chosen so that there is an easy mapping between WP3 sections in the Description of Work (DOW) and the sections of the Deliverable. For each technology, an overview of relevant state-of-the-art and general uses is provided, along with appropriate references. Moreover, the application of the technology in relationship with the CrossCult goals is explained, considering the requirements elicited for the four pilots in Deliverable D2.1 ("*Pilot specifications*"). Finally, each technology is accompanied by a table with software that can be used directly or indirectly in the project (considering functional requirements, non-functional ones and licenses) or simply act as inspiration or reference for the expected outcomes of WP3.

## 1.3.   Structure of the deliverable

The document follows the structure of the CrossCult WP3, where five distinct tasks about technologies are identified, accompanied by a sixth task on integration and testing. The sections of the deliverable, therefore, are:

- **Section 2: User modelling, recommendation and personalisation.** This section looks at the task of capturing, representing and managing information about individuals and groups, and with the implementation of the modules in charge of matching the information available about users, venues and context in order to deliver the most meaningful experiences and content.

- **Section 3: Machine learning, semantic reasoning and crowdsourcing.** This section considers the mechanisms that may be used to consolidate and enhance the CrossCult knowledge bases through the computational discovery of unknown associations between characters, events, cultural assets and venues within the same or different digital repositories.

- **Section 4: Context mining and processing.** This section focuses on the identification and processing of pieces of contextual information that may be used to maximise the value of the information delivered to the users.

- **Section 5: Sporadic social networks and crowd management.** This section deals with the implementation of mechanisms to promote socialisation, group making and communication, while managing the crowds in the venues so as to avoid congestion points and ensure good compromise between the users' expectations and the time required to see and interpret the items.

- **Section 6: Visualisation of associations and micro-augmentations.** This section focuses on the innovations related to the user's visualisation of the associations among historical facts and pieces of cultural heritage, reinforced with small stimuli to foster awareness and to maintain user engagement.

- **Section 7: Integration and alpha testing.** Finally, this section looks at the tools and methodologies that may be applied to test the developed modules in controlled environments, prior to deploying the solutions in the venues and pilots of WP5.

# 2. User modelling, recommendation and personalisation

## 2.1.  Profiling

### 2.1.1. Survey

Personalisation in cultural heritage places requires getting information about each visitor, i.e. establishing a profile. This remains a big issue, mainly because of the short duration of most visits (Falk et al., 2010) and the fact that most visitors might visit a specific place only once. Within these time and place restrictions, visitor profiles need to be created quickly and effectively in terms of their appropriateness for the different visitors. The question raised though is how to start creating these profiles and where to find the necessary information. More than in virtual spaces, where information can be more easily obtained, personalisation applications in physical spaces face the well-known *cold start* problem (Kuflik et al, 2011). Furthermore, implicit profiling techniques are preferred, because users may not be willing to directly enter information through profile completion pages since the perceived risks are high but the benefits are not always apparent (Kelly and Teevan, 2003) and they can move them away from the visit goal itself.

Different approaches have been employed like using heuristic techniques (Hameed et al., 2012), probability-based algorithms (Lin et al., 2014) or user prototypes/personas (Roussou et al., 2013). The most common approaches to gather visitors' interests use questionnaires or short quizzes (Vayanou et al., 2014). However, visitors are often reluctant at form-filling activities (Aoidh et al., 2012), making it necessary to apply indirect approaches for the creation of user profiles. For example, the PIL project gathers information from users' interaction with a specially designed webpage and uses this information to create personal profiles prior to the actual visit (Kuflik et al., 2011). Other approaches exploit user stereotypes (Roussou et al., 2013), using personas to augment visitor experience with personalised mobile storytelling. Yiakoumettis et al. (2014) recommend optimised routes based on personal interests computed with a relevance feedback approach using preferences given by users on a set of object pairs. Another original approach can be found in Fosh et al. (2014), who suggest recommendations based on visitor interpretations especially tailored for one of their friends or loved ones. Bohnert et al. (2012) and Semeraro et al. (2012) tried to predict user behaviour during the visit to drive recommendations, based on previous user feedbacks (e.g. ratings, tags). Bonis et al. (2009) asked visitors to customise avatars, since it is known that customised avatars reflect personality (Fong et al., 2015). Antoniou and Lepouras (2010) observed visitor movements within museums and inferred personality profiles. Recently, it has been repeatedly suggested that mini-games personality traits can successfully solve the cold start problem (Tkalcic and Chen, 2015; Antoniou et al., 2013a; Antoniou et al., 2013b; Naudet et al. 2013; Lykourentzou et al., 2013). Generally, implicit profiling approaches are necessary to confirm and refine explicit profiles. The observation of the visitor behaviour can reveal interests. For example, a visitor interested by a certain type of exhibits will spend more time in front of these exhibits while making less frequent stops for less appreciated exhibits (Castagnos and Boyer, 2006; Bohnert et al., 2007). Additionally, other aspects can be revealed, like, e.g. the visiting style (Véron et Levasseur, 1989), or the visitor's fatigue (Davey, 2005), which are of particular interest.

In the last few years, a noticeable research trend has focused on the mining of social networks to build user profiles. It has been demonstrated that people's behaviour on such media does indeed reflect real personality preferences and real behaviour (Back et al., 2010; Gosling et al., 2011). Therefore, social media can be rich sources of necessary information in order to efficiently create user profiles for personalised content in cultural heritage, also solving the cold start problem, as Tkalcic and Chen (2015) suggest. Different social media have been used for the elicitation of user profiles like YouTube and Pinterest (Feinberg et al., 2012), Instagram (Weilenmann et al., 2013), Twitter (Lin et al., 2012) and Facebook (Talavera et al., 2012). For example, it was found that extraversion could be successfully calculated from the frequency of Facebook use, the high activity, the number of friends, number of attended events, and the intensity of photo postings (Bhargava et al., 2015; Gosling et al., 2011; Ortigosa et al., 2011). The Prediction API developed at the Psychometrics Centre, University of Cambridge (http://applymagicsauce.com/you.html) analyses Facebook data to calculate basic personality elements. Celli et al. (2014) showed that the analysis of profile photos in social media can also provide valuable personality information. However, as shown by Theodoridis et al. (2015), research is still at the early stages and algorithms are not always reliable. Furthermore, all the aforementioned pieces of work may apply for individuals and the elicitation of individual user profiles, while not much has been done in the field of adaptation for group visitors and elicitation of group profiles in cultural heritage, despite the fact visitors rarely visit alone (Petrelli et al., 2005).

## 2.1.2. CrossCult applications

The uses of profiling technologies in CrossCult relates to the capture, representation and management of information about individuals and groups. Quick profiling strategies will be developed or adapted from previous works, including carousels of predefined items, quizzes and games, and exploiting information from social networks. Likewise, different feedback mechanisms will be implemented to progressively update the profiles, including explicit forms of feedback (e.g. likes/dislikes, ratings and comments) as well as implicit ones (e.g. the fact that a user stands in front of an item, the number of times she replays an audio/video clip, or whether he/she shares a piece of content with others). Implicit profiling and continuous user profile update will be mainly based on a visitor behavioural model (see e.g., (Osche and al., 2016) for a preliminary proposition) and usage mining techniques (Srivastava et al., 2000). Different techniques like multiple linear regression or graphical modelling will be used to analyse the observations of the visitors' behaviour, focusing on maximising a global satisfaction function. Last, deep learning approaches will be experimented and might reveal new aspects of the user model or unexpected correlations.

## 2.1.3. Software, tools and algorithms

**Table 1: Profiling software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|---|---|---|---|
| **Apache Spark** | Apache License Version 2.0 | Apache Spark is a fast and general-purpose cluster computing system. It provides high-level APIs in Java, Scala, Python and R and an optimised engine that supports general execution graphs. It also supports a rich set of higher-level tools and | Use it as a framework to support the building of computationally extensive tasks such as background analysis of user profiles to |

| | | structured data processing, MLlib for machine learning, GraphX for graph processing, and Spark Streaming. | update recommendations. |
|---|---|---|---|
| **Prediction API** | Non-exclusive commercial license | Apply Magic Sauce (AMS) is a service that predicts users psycho-demographic traits based on their digital footprints. Currently, the service accepts Facebook Like IDs as the input data format, and makes predictions available for personality, satisfaction with life, intelligence, age, gender, sexual preference, interest in given area, political views, religion, relationship status, and leadership potential. | Use as a source of personality data for quick profiling. |
| **MyPersonality database** | Non-commercial license, only for academic research that does not earn revenue, not in collaboration with any commercial entities | myPersonality was a popular Facebook application that allowed users to take psychometric tests. As well as the data from the tests, around 40% of the respondents agreed to give access to their Facebook profile data, and social network data. http://mypersonality.org/ makes available some of the information stored in the myPersonality databases to the wider academic community. Currently, the database contains more than 6,000,000 test results together with more than 4,000,000 individual profiles. Respondents come from various age groups, backgrounds and cultures. Their scores are combined with large amounts of additional information from those who opted in to sharing it, including detailed demographic profiles, records of behaviour in Facebook, interests, preferences, opinions, etc. | Use as training data for profiling algorithms. |
| **Discover my profile** | Non-commercial license | This is a 2015-relaunch of the myPersonality Facebook app, which currently allows users to take 27 tests: IQ score, personality profile based on the classical "big 5" personality traits, etc. It also allows affiliated researchers and students to host their own assessments and offer feedback to up to 35,000 daily participants around the world. | Use as a means to get personality and cognitive profile of a user. Possibly a place to host CrossCult tests and get input from an existing large community of visitors. |
| **International Cognitive Ability Resource (ICAR)** | Under development | ICAR is a public-domain and open-source tool that will provide a large and dynamic bank of cognitive ability measures for use in a wide variety of applications. It builds on automatic item generation techniques that yield items with predictable psychometric qualities. These will be distributed as functions in psychometrically-informed, easily-implemented, open-source software. Suitable statistical methodologies will be developed and sub-projects will use ICAR items to explore specific research questions. | Use as a means to get the cognitive profile of a user. |
| **PredictionIO** | Apache License, Version 2.0 | PredictionIO is an open-source machine learning server for software developers to create predictive features, such as personalisation, recommendation and content discovery. | Use to build a production-grade engine to predict users' preferences and personalise content. |
| **MyMuseumStory** | UOP proprietary | Facebook game allowing extracting the cognitive style of the user. | Use as a means to get the cognitive profile of a user. |
| **Ventour & Ventour APPs** | GVAM proprietary | App builder and content management system for guiding apps in museums. Direct profiling tools to configure visits and personaize narrations and interaction with games and modules | Platform to gather and publish direct profile information from app users |
| **Deeplearning 4j** | Apache 2 Licence | Deep learning library | Implicit user profiling. |

| (http://deeple arning4j.org/) | | | |
|---|---|---|---|

## 2.1.4. References

Abdul Hameed, M., Ramachandram, S. and Al Jadaan, O. 2012. Information gain clustering through prototype-embedded genetic K-mean algorithm (IGCPGKA): A novel heuristic approach for personalisation of cold start problem. In Proceedings of 2nd International Conference on Computational Science, Engineering and Information Technology, pp. 390-395.

Antoniou A. and Lepouras, G. 2010. Modeling visitors' profiles: A study to investigate adaptation aspects for museum learning technologies. Journal on Computing and Cultural Heritage 3(2).

Antoniou, A., Lykourentzou, I., Rompa, J., Tobias, E., Lepouras, G., Vassilakis, C. and Naudet, Y. 2013a. User profiling: Towards a Facebook game that reveals cognitive style. In Games and Learning Alliance, pp. 349-353, Springer International Publishing.

Antoniou, A., Lepouras, G., Lykourentzou, I., Naudet, Y. 2013b. Connecting physical space, human personalities, and social networks: The Experimedia Blue project. In Proceedings of International Biennial Conference Hybrid City, Subtle Revolutions.

Aoidh, E.M., Bertolotto, M. and Wilson, D. 2012. Towards dynamic behavior-based profiling for reducing spatial information overload in map browsing activity. GeoInformatica 16: 409–434.

Back, M.D., Stopfer, J.M., Vazire, S., Gaddis, S., Schmukle, S.C., Egloff, B. and Gosling, S.D. 2010. Facebook profiles reflect actual personality, not self-idealization. Psychological Science 21(3): 372–374.

Bhargava, P., Brdiczka, O. and Roberts, M, 2015. Unsupervised modeling of users' interests from their Facebook profiles and activities. In Proceedings of 20th International Conference on Intelligent User Interfaces, pp. 191–201.

Bohnert, F., and Zukerman, I. 2007. Using viewing time for theme prediction in cultural heritage spaces. In AI 2007: Advances in Artificial Intelligence, pp. 367–376, Springer.

Bohnert, F., Zukerman, I., Laures, J. 2012. GECKOmmender: Personalised theme and tour recommendations for museums. In Proceedings of International Conference on User Modeling, Adaptation and Personalization, pp. 26-37.

Castagnos, S. and Boyer, A. 2006. A client/server user-based collaborative filtering algorithm: Model and implementation. In Proceedings of 17th European Conference on Artificial Intelligence, pp. 617–621.

G. Davey G., 2005. What is museum fatigue? Visitor Studies Today, 8(3): 17–21.

Falk, J.H., Koran Jr., J.J., Dierking, L.D. and Dreblow, L. 2010. Predicting visitor behavior. Curator: The Museum Journal 28: 249–258.

Feinberg, M., Geisler, G., Whitworth, E. and Clark, E. 2012. Understanding personal digital collections: An interdisciplinary exploration. In Proceedings of Designing Interactive Systems Conference, pp. 200-209.

Fosh, L., Benford, S., Reeves, S., Koleva, B. and Brundell, P. 2013. See me, feel me, touch me, hear me: Trajectories and interpretation in a sculpture garden. In Proceedings of SIGCHI Conference on Human Factors in Computing Systems, pp. 149-158.

Gosling, S., Augustine, A.A., Vazire, S., Holtzman, N. and Gaddis, S. 2011. Manifestations of personality in online social networks: Self-reported Facebook-related behaviors and observable profile information. Cyberpsychology, Behavior and Social Networking 14(9): 483-488.

Kelly, D. and Teevan, J. 2003. Implicit feedback for inferring user preference: A bibliography. ACM SIGIR Forum 37(2): 18-28.

Kuflik, T., Sheidin, J., Jbara, S., Goren-Bar, D., Soffer, P., Stock, O. and Zancanaro, M. 2007. Supporting small groups in the museum by context-aware communication services. In Proceedings of 12th international Conference on Intelligent user interfaces, pp. 305-308.

Lin, C., Xie, R., Li, L., Huang, Z. and Li, T. 2012. PRemiSE: Personalized news recommendation via implicit social experts. In Proceedings of 21st ACM International Conference on Information and Knowledge Management, pp. 1607-1611.

Lykourentzou, I., Claude, X., Naudet, Y., Tobias, E., Antoniou, A., Lepouras, G., Vassilakis, C. 2013. Improving museum visitors' Quality of Experience through intelligent recommendations: A visiting style-based approach. In Intelligent Environments (Workshops) pp. 507-518.

Naudet, Y., Lykourentzou, I., Tobias, E., Antoniou, A., Rompa, J. and Lepouras, G. 2013. Gaming and cognitive profiles for recommendations in museums. In Proceedings of 8th International Workshop on Semantic Media Adaptation and Personalization, pp. 67-72

Ortigosa, A., Quiroga, J.I. and Carro, R.M. 2011. Inferring user personality in social networks: A case study in Facebook. In Proceedings of 11th International Conference on Intelligent Systems Design and Applications, pp. 563-568.

Osche, P.-E., Castagnos, S., Napoli, A. and Naudet, Y. 2016. Walk the line: towards an efficient user model for recommendations in museums. In Proceedings of 11th International Workshop on Semantic and Social Media Adaptation and Personalization.

Petrelli, D. and Not, E. 2005. User-centred design of flexible hypermedia for a mobile guide: Reflections on the HyperAudio experience. User Modeling and User-Adapted Interaction 15(3-4): 303-338.

Roussou, M., Katifori, A., Pujol, L., Vayanou, M. and Rennick-Egglestone, S.J. 2013. A life of their own: Museum visitor personas penetrating the design lifecycle of a mobile experience. In Proceedings of CHI Conference on Human Factors in Computing Systems, Extended Abstracts, pp. 547-552.

Santana Talavera, A., Rodríguez Darias, A.J., Díaz Rodríguez, P. and Aguilera Ávila, L. 2012. Facebook, heritage and tourism reorientation. The cases of Tenerife and Fuerteventura (Canary Isles, Spain). International Journal of Web Based Communities 8(1): 24-39.

Semeraro, G., Lops, P., de Gemmis, M., Musto, C. and Narducci, F. A folksonomy-based recommender system for personalized access to digital artworks. ACM Journal on Computing and Cultural Heritage 5(3).

Srivastava, J., Cooley, R., Deshpande, M. and Tan, P.-N. 2000. Web usage mining: Discovery and applications of usage patterns from web data. ACM SIGKDD Explorations Newsletter, 1(2): 12–23.

Theodoridis, T., Papadopoulos, S. and Kompatsiaris, Y. 2015. Assessing the reliability of Facebook user profiling. In Proceedings of 24th International Conference on World Wide Web, pp. 129-130

Tkalcic, M. and Chen, L. 2015. Personality and recommender systems. In Recommender Systems Handbook, pp. 715-739. Springer US.

Vayanou, M., Katifori, V., Chrysanthi, A. and Antoniou, A. 2016. Collocated interaction in cultural storytelling experiences: How to coordinate visitor actions? In Proceedings of 19th ACM Conference on Computer-Supported Cooperative Work and Social Computing.

Véron, E. and Levasseur, M. 1989. Ethnographie de l'exposition : l'espace, le corps et le sens. Centre Georges Pompidou, Bibliothèque publique d'information.

Weilenmann, A., Hillman, T. and Jungselius, B. 2013. Instagram at the museum: Communicating the museum experience through social photo sharing. In Proceedings of SIGCHI Conference on Human Factors in Computing Systems, pp. 1843-1852.

Yiakoumettis, C., Doulamis, N.D., Miaoulis, G. and Ghazanfarpour, D. 2014. Active learning of user's preferences estimation towards a personalized 3D navigation of geo-referenced scenes. GeoInformatica 18(1): 27-62.

## 2.2. Accessibility

### 2.2.1. Survey

Various local and international laws provide accessibility guidelines and regulation for specific activities or events to be adapted to people with disabilities. They also provide tools or devices that permit access to the content autonomously.  As Benoît Villain (Museum LaM of Lille) said, "*mobile devices contribute to the contemporary appearance of the museum*", but their great contribution is to be able to ensure maximum accessibility of content in equivalent condition to the general public, and have the ability to promote audience participation.

Depending on the conditions and requirements of different users, there is an extensive list of accessibility resources and several software options to produce and integrate these resources in the mobile devices, audio guides, sign guides, etc. The question is what the main approaches are that international museums use to guarantee access using new technologies. It may be said that two main approaches coexist:

- **Specific designs:** some international museums have adopted accessibility measures fragmenting their audiences with different guides, according to their needs. For example, the British Museum offers different products depending on the abilities, necessities or preferences of their publics: general audio guide, family guide, sign guide and audio descriptive guide.

- **Universal Design:** accessibility measures should work in a logical and consistent manner, be compatible with each other and configurable from a single application. That is, the amplitude of the catalogue of measures should not lead to the diversification of them in different systems and media. The "*Design for all*" concept promotes equal opportunities for access and use of technologies and services, development tools and contents that are capable of adapting to every visitor, regardless of their abilities, age, gender, language or cultural background.

TATE Modern was the first cultural institution that included a sign guide in its offer. This museum represents well the evolution of the accessibility in this context. Its multimedia guide offers several tours adapted to the necessities of people with disabilities (e.g. the "*British Sign Language tour*" offered by the National Gallery). The same model has been elected by GVAM in several museums in Spain (e.g. National Archaeological Museum or National Heritage Palaces) (Pajares & Solano, 2012). The principal goal of this approach is to achieve a unique tool for all people with an adaptive design and specific tour or contents. The perfect answer to this objective is in the mobile device of each visitor. Other international institutions, like Guggenheim, are working in this way.

## 2.2.2. CrossCult applications

To personalise and to make a service accessible are not the same thing. Although the accessible functions and content are part of the personalisation service, there is something that distinguishes them significantly: while personalisation is fundamental, accessibility is also a right of persons with disabilities, without which the project and its applications could breach various regulations and conventions at a national and international level. The adaptation to this requirement of the cultural institutions means they should introduce accessibility criteria in all its spaces, routes and means, and must comply with technical regulations. To guarantee those rights, museums must offer clear information on specific activities or events adapted to people with disabilities and tools or devices which permit access to their content autonomously.

Prioritising "*for all*" applications in CrossCult should manage in a correct way the accessibility resources listed below to be considered an accessible application "*for each one*". Prioritising appropriately the accessibility resources below, the CrossCult applications should allow the user to enable or disable the next options depending on his/her needs:

- **Closed captioning** is a "*service to support communication on the screen, with text and graphics, oral discourse, suprasegmental information and sound effects that occur in any audiovisual work*" (UNE 153010). More than 90% of people with hearing disabilities will access the information in our guide through the subtitles.

- **Sign language videos:** While those deaf people who can read lips choose to activate the subtitles, the signers prefer to activate the sign language. Some deaf people are bilingual, as they can read lips as well as understand sign language –in this case it is possible to activate both resources. A sign language (also signed language) is a language that, instead of acoustically conveyed sound patterns, uses manual communication and body language

to convey meaning. In the context of a media guide, it is used to display a video with a sign language interpreter who translates text content and/or the audio that is played on each screen. We must remember that there are different sign languages (or signs), one for each of the major world languages and an international sign language. The latter is an artificial mixture of the signs of the different languages that have formed a new one.

- **Verbal description guides or audio-description for blind people or visually impaired users:** Audio-description involves the interpretation of a visual scene for verbal re-explanation and description to people who cannot see or have difficulties in doing so. It is a resource aimed primarily at people with visual disabilities, but is also an important reinforcement for users with intellectual or sensory disabilities. The techniques applied in the audio description of films differ significantly from those applied in the audio description of exhibitions, although the criteria to be followed in both cases can be found in the Spanish standard UNE 153020.

- **Audio-navigation or integration/support with screen reader options:** This is a form of guided navigation through the audio description, which allows visually impaired users to navigate independently thorough the available options, menus or content of an interactive application.

- **Pictograms:** The pictogram is an educational resource that facilitates communication between people with cognitive or intellectual disabilities, and those with insufficient knowledge of the language. They can be icons, signs or schematic images associated with activities, situations, environments or actions. A good example of its application is in the pictogram version of the UN Convention on the Rights of Persons with Disabilities, which is edited by the CEAPAT*2*.

- **Easy reading:** Easy reading consists in simplifying texts in order to make information and culture accessible to all people with reading difficulties. Different legal texts, both national and international, recommend that public and private entities broadcast their messages in an easy to comprehend language. Easy reading means adapting text so that it is understandable, while using a varied and adult vocabulary.

- The **magnifier** is a tool that visually impaired visitors can use to increase the size of screen elements that are not perceived clearly.

- A **high contrast** option will put all the white letters on black backgrounds, helping many people with residual vision to perceive more easily the text on display.

- An **automatic sweep** will make activating each display option easier for people with mobility problems in their arms or hands. Automatic scanning tabulates sequentially for each option available on the screen pausing a few seconds for each one. The desired option can be activated by pressing just one button when this is selected by the system, minimising movement and the precision required to access all the contents of the guide.

## 2.2.3. Software, tools and algorithms

**Table 2: Accessibility software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|---|---|---|---|
| **Voice Over** | SaaS | VoiceOver allows the user to access their Macintosh/iOS device based on spoken descriptions and the keyboard. Developers can use WAI AREA directives to comply and use it for IOS APPs | The feature can be used as inspiration for how to increase accessibility for blind and low-vision users, as well as for users with dyslexia. |
| **Google TalkBack** | SaaS | TalkBack increases accessibility via spoken, audible, and vibration feedback. Developers can use WAI AREA directives to comply and use it for Android 4.x APPs | The feature can be used as inspiration for ways to help blind and vision-impaired users interact with their devices. |
| **Vocalia** | SaaS | Vocalia is a voice control platform based on natural languages technologies that allows visual impaired people to control complex interfaces and access knowledge bases | Apps control by natural language and voice recognition for blind people and people with functional disabilities |
| **Magic** | Proprietary | MAGic is a screen magnification and screen reading solution for low vision computer users. It delivers smooth, crisp letters, even at the highest magnification levels. MAGic delivers smooth, crisp letters, even at the highest magnification levels. | The feature can be used as inspiration for ways to help visually impaired users. |
| **Zoomtext Xtra** | Proprietary | ZoomText is a magnification and screen reading software for the visually impaired, helping countless individuals stay employed and in touch with the world. | The feature can be used as inspiration for ways to help visually impaired users. |
| **Supernova** | Proprietary | Screen magnifier and screen reader | The feature can be used as inspiration for ways to help visually impaired users. |
| **Colour Contrast Analyzer** | Proprietary | This extension allows you to analyse text colour contrast problems on a webpage according to the WCAG 2 text colour contrast requirements. | This feature can be used to test applications of CrossCult (especially web-based ones) for accessibility requirements. |
| **CSS Analyzer** | Proprietary | This service has been provided to allow you to check the validity of your CSS against the W3C's validation service, along with a colour contrast test, and a test to ensure that relevant sizes are specified in relative units of measurement. | This feature can be used to test applications of CrossCult (especially web-based ones) for accessibility requirements. |

## 2.2.4. References

British Museum audio guides.
http://www.britishmuseum.org/visiting/planning_your_visit/audio_guides.aspx?lang=en

National Gallery, 2016. British sign language tours. https://www.nationalgallery.org.uk/visiting/audio-tours/british-sign-language-tour

Pajares, J.L. and Solano, J. 2012. Museums of the future. The role of accessibility and mobile technologies. White Paper. http://www.gvam.es/ebook/ebook_MuseosDelFuturo.pdf

United Nations, 2006. Convention on the rights of persons with disabilities.
http://www.un.org/disabilities/convention/conventionfull.shtml

UNE 153010 norm. 2012. Subtitling for deaf and hard-of-hearing people.
http://www.aenor.es/aenor/normas/normas/fichanorma.asp?tipo=N&codigo=N0049426#.V-5NiJOLQUF

UNE 153020 norm. 2005. Audio description for visually impaired people. Guidelines for audio
description procedures and for the preparation of audio guides.
http://www.aenor.es/aenor/normas/normas/fichanorma.asp?tipo=N&codigo=N0032787&PDF=Si#.V-5N6JOLQUE

Villain, B. 2011. CLIC France Interview. http://www.club-innovation-culture.fr/interview-benoit-villain-expose-les-dispositifs-mobiles-du-lam-a-lilles/

## 2.3.  Recommendation

### 2.3.1. Survey

Personalised recommendation is about suggesting things to the user, in a way that takes into account their preferences or interests, sometimes their situation or the context affecting them. In the last twenty years, different categories of recommender systems exploring many different algorithms have been proposed (see, e.g., the Recommender Systems Handbook, Ricci et al., 2015, for a survey). Approaches that have been selected for the CrossCult context include classical numerical approaches, which mainly include collaborative filtering techniques and diversity-based algorithms, in addition to content-based filtering, and approaches exploiting formalised knowledge (e.g. ontologies).

While classical recommender systems operate on two elements (namely, users and items), *Context-Aware Recommender Systems* (CARS) consider context as a third element (Adomavicius et al., 2005). This is particularly pertinent in a mobile context, where the user's situation can change and consequently the pertinence or validity of user's needs, interests or preferences. With the rise of mobile devices, recommender systems are now used in physical spaces to provide location-aware recommendations to users. But apart from the user's location, other context variables need to be taken into account, like, e.g. time constraints, social context (e.g. visiting in group: Martin and Trummer, 2005; Liiv et al., 2009), or the importance of exhibits for the museum. In all cases, personalisation in single-use systems like museums seems to accept a trade-off between time efficiency (given the fact that most museum visits have particularly short durations, Falk et al., 1985) and accuracy of user profiles (Noor et al., 2009). In the domain of tourism and cultural heritage, over the past years, guiding applications exploiting geolocation technologies and including recommendation features have been proposed. Applied to museums in particular, recommender systems in physical spaces take different forms: recommendation of exhibits to see (e.g. Bohnert et al., 2012), route/path recommendation or path optimisation (e.g. Mathias et al., 2014; Basile et al., 2009), personalised multimedia content proposed in relation to the exhibition (e.g. Martin and Trummer, 2005), and adaptation of the mediation (e.g. Zhou et al., 2009,

Gavalas and Kenteris, 2011). Some approaches propose a mix, like in (Liiv et al., 2009) or (Fournier and Viennet, 2015).

*Collaborative filtering* techniques aim at identifying the active user with a set of persons having the same tastes, based on their preferences and past actions (Castagnos et al., 2008). Typically, those algorithms build a numerical user profile modelling interests from provided ratings, aggregating all in a user-item rating matrix. Different approaches allow exploiting the rating matrix to adapt recommendations according to similar users' preferences, e.g., clustering techniques (Castagnos et al., 2006), matrix factorisation (Koren, 2009), neighbourhood-based models (Brun et al., 2011), Bayesian networks (Castagnos et al., 2005), or neural networks (Hidasi et al., 2016).

*Diversity-aware recommenders* focus on introducing diversity in recommendations, rather than only focusing on similarity and precision. In general, it has the potential to significantly enhance the efficiency of recommendations (McGinty and Smyth, 2003). Traditional methods for diversifying search results rely on attribute-based diversification (Bradley and Smyth, 2001) or re-ranking of the top-N search results (Teevan et al., 2005), based, e.g., on results homogeneity reduction (Anagnostopoulos et al., 2005) or on the analysis of query reformulations (Radlinski and Dumais, 2009). More recent approaches try to maximise diversity while maintaining adequate similarity with a binary optimisation problem (Zhang and Hurley, 2008); to evaluate novelty focusing on diverging from popular items-only recommendations (Celma and Cano, 2008); to predict a user's unknown items (Hijikata et al., 2009), focus on enhancing the users' perception and acceptance rate while introducing diversity (Castagnos et al, 2013). Diversity can not only be exploited for recommendations re-ordering, enrichment or filtering, but also for offering topic diversification (Ziegler et al., 2005), which has proven to improve user satisfaction. Last, because users' need for diversity can change over time (Castagnos et al., 2010), approaches have been proposed to monitor this need (L'Huillier et al., 2014) and adapting recommendations based on detected contexts (L'Huillier et al., 2016).

Finally, the exploitation of formalised knowledge has led to two related categories of recommender systems (RS): *knowledge-based RS* (Burke, 2000) and semantic, ontology-based, or *semantics-aware RS* (de Gemmis et al., 2015). All have in common that recommendations are computed using explicit domain knowledge about the items, the users, potentially the context, and their relationships. The algorithms compute a utility score for each item, representing a prediction of the user's interest for this item, meaning its fitting with the user interests, preferences or needs, modelled in a user profile. Knowledge-based recommender systems rely on knowledge formalising how items characteristics relate to a user interests, preferences or needs (Ricci et al., 2015), while semantics-aware ones exploit the semantics of concepts formalised in knowledge bases, and especially using ontologies, e.g. via collaborative filtering (see e.g. Ziegler et al., 2004; Mobasher et al., 2004) and content-based filtering (Jannach et al., 2010, de Gemmis et al., 2015).

## 2.3.2. CrossCult applications

Many of the features promised by the CrossCult platform depend on the implementation of modules capable of matching the information available about users, venues and context in order

to deliver the most meaningful experiences and content. The goal is to make users gain time by automatically suggesting items in accordance with their preferences and behaviours (recommendation), or by pulling items that fit their queries in search engines (personalisation). The project will first look at the previous work on building a personalised guide for museums, named MyMuseumGuide (Naudet et al. 2015). The system integrates a recommender system exploiting in particular the user cognitive style, in addition to their interests. In this context the authors define a visitor model, defined as a tuple or ordered list, including a cognitive profile, personal interests, personal profile, location, activity and time constraints. The recommender itself computes the sequence of exhibits that best matches each visitor, including actions that are suggested.

Sequences/paths recommendations will be addressed, with the aim of (i) maximising the overall satisfaction rate and the value created with the recommendations, while (ii) integrating the historic constraints (e.g., screenwriting, reinterpretation of history, scenery) and (iii) improving the visitor traffic (e.g. reducing congestion points). The approach will integrate mainly the hybridisation of different recommender algorithms, and in particular numerical collaborative and diversity-based filtering approaches with semantics-aware approaches, the rearrangement of recommendations according to human factors (e.g., visiting style, fatigue, preferences, tolerance to distance, technological intrusion) and on the museum characteristics (e.g. crowd level, congestion).

In CrossCult, the semantics-aware recommendations will be based on previous work from (Naudet et al. 2014), adapted to mobile location-aware context. In the related recommender system prototype named Sphynx, the recommendation algorithm exploits filtering and reasoning to build an inferred knowledge base on which matchmaking between user and items is performed. A semantic similarity measure combined with category matching provides a score for each item according to users' interests and context.

Recommender systems in CrossCult will also exploit trending topics information from social media to allow venue content curators to promote exhibits or collections relevant to topical issues; this feature will enable venues to take advantage of subjects' publicity, turning it into interest on their content. Trending topics are read from social media and matched against the metadata of exhibits and collections; subsequently, the highest matches are presented to venue content curators who can then choose how (and if) these matches can be best used to trigger visitor interest and promote the venue content. Some cases will be developed of recommenders using social media info from the user to make personalised recommendations according with his/her social media profile or interactions.

Regarding the aspect of group adaptation, CrossCult will create group profiles and adapt content according to the group's requirements. To this end, the personalisation mechanisms in the CrossCult project include algorithms to determine whether an individual is visiting alone, or in a group. While this can be extracted using a direct question, additional technological means can be exploited to facilitate user profiling: near-field communications allows users of modern smartphones and tablets to initiate a relationship (Coskun et al., 2015). Similarly, identifying that users that have an established "friend" relationship in social networks are visiting a cultural space

at the same time, indicates that the relevant users are visiting as a group (or have an opportunity to form one, if the collocation is coincidental).

Finally, recommenders in CrossCult will foster the creation of sporadic social networks, by recommending connections among people that co-visit the venue. In this respect, the profiles of current venue visitors will be analysed and groups of people with related interests will be determined; then suggestions will be forwarded to members of each group to join an online group through which they can engage in collaborative live discussion, participate in educational games or even formulate a physical group. These recommendations will respect user anonymity, and only after explicit consent of all interested parties the electronic IDs of the users will be exchanged. We expect that user interaction within the context of sporadic social networks will promote the creation of new views on the venue content and the connections between exhibits, enriching the venue's informational content. Some of the connections are expected to outlive the duration of the visit (e.g. through an establishment of a "friend" relation in social networks), supporting post-visit actions (including discussions, diffusion of information to other social network members and so forth), which could extend the benefits of the visiting experience.

## 2.3.3. Software, tools and algorithms

**Table 3: Recommendation software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|---|---|---|---|
| **MyMuseumGuide** | LIST proprietary | A mobile application including in particular a recommending engine computing sequences of POIs for a museum visit, exploiting both user interests and its cognitive profile. | The recommendation engine can be used as a basis for recommendations of POIs and Paths in museum. To be extended with semantic algorithms. |
| **Sphynx** | LIST proprietary | Software in Java implementing an ontology-based context-aware recommendation system | To be reused as a basis for a semantics-aware recommender. |
| **LensKit** | GNU Lesser General Public license version 2.1 | LensKit is free and open-source software that provides building blocks to create recommender systems for a range of applications. A LensKit recommender comprises a set of interfaces providing recommendation, rating prediction, and other related services using one or more recommender algorithms connected to a data source. | Use as a conceptual framework for the construction of a generic recommender system; possibly reuse some constructs and algorithms. |
| **Rdf4j** | Eclipse Distribution License - v 1.0 | RDF4J (formerly Sesame) is an open source Java framework for processing RDF data. This includes parsing, storing, inferencing and querying of/over such data. It offers an easy-to-use API that can be connected to all leading RDF storage solutions. It allows connecting with SPARQL endpoints and creating applications that leverage the power of linked data and Semantic Web. | Use as a Java framework to query /manage a semantic knowledge base via a SPARQL endpoint. |
| **Apache Jena** | Apache License Version 2.0 | A free and open source Java framework for building Semantic Web and Linked Data applications. | Use as a Java Framework to query /manage a semantic knowledge base via a SPARQL endpoint. |
| **Apache** | Apache | The Apache Mahout™ project's goal is to | Use it as a framework for |

| Mahout | License Version 2.0 | build an environment for quickly creating scalable performance machine learning applications. | complex algorithm implementation. |
|---|---|---|---|
| **Virtuoso** | Commercial | The knowledge base engine that powers DBpedia. | To be used a semantic knowledge base for reasoning. |
| **Ontotext GraphDB** | Commercial | A semantic graph database that serves organisations to store, organize and manage content in the form of semantically enriched smart data. | To be used as a semantic knowledge base for reasoning. |
| **SemSim** | Apache 2 | Semantic similarity library written in Java. | Used in the recommendation process to compute the similarity between different items (e.g. POIs). |
| **Elastic search** | Apache 2 | A document oriented NoSQL DB engine that features fast indexing, powerful full-text search capabilities and that is built to run on distributed systems. | Use for fast indexing and basic recommendations based on search relevance score. |
| **LibRec** (http://www.librec.net) | GNU GPL | Library of recommender algorithms | Numeric approaches to use for hybridisation. |
| **CARSKit** (https://github.com/irecsys/CARSKit) | GNU GPL | Library for context-aware recommender systems | CARS for hybridisation. |
| **DANCE** | Property of LORIA/CNRS | Diversity-based model | Use for hybridisation. |
| **Sifarish** https://github.com/pranab/sifarish) | Open-source | Content based and collaborative filtering based recommendation and personalization engine implementation on Hadoop and Storm | Use for hybridisation. |

## 2.3.4. References

Adomavicius, G., Sankaranarayanan, R., Sen, S., and Tuzhilin, A. Incorporating contextual information in recommender systems using a multidimensional approach, ACM Transactions on Information Systems, 23(1): 103–145.

Anagnostopoulos, A., Broder, A.Z. and Carmel, D. 2005. Sampling search-engine results. In Proceedings of 14th International Conference on World Wide Web.

Basile, P., de Gemmis, M., Iaquinta, L., Lops, P., Musto, C., Narducci, F., & Semeraro, G. 2009. SpIteR: A module for recommending dynamic personalized museum tours. In Proceedings of IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology, pp. 584-587.

Bohnert, F., Zukerman, I. and Albrecht, D.W. 2012. Realistic simulation of museum visitors' movements as a tool for assessing sensor-based user models. In Proceedings of 20th International Conference on User Modeling, Adaptation and Personalization, pp. 14-15.

Bradley, K. and Smyth, B. 2001. Improving recommendation diversity. In Proceedings of 12th Irish Conference on Artificial Intelligence and Cognitive Science.

Brun, A., Castagnos, S. and Boyer, A. 2011. From community detection to mentor selection in rating-free collaborative filtering. Advances in Multimedia.

Burke, R. 2000. Knowledge-based recommender systems. In Encyclopedia of Library and Information Systems. Marcel Dekker.

Castagnos, S., Boyer, A. and Charpillet, F. 2005. Vers un filtrage collaboratif distribué : le modèle RSB. 3èmes journées francophones sur les Modèles Formels de l'Interaction (MFI 2005), Caen, France, Mai 2005.

Castagnos, S. and Boyer, A. 2006. A client/server user-based collaborative filtering algorithm: Model and implementation. In Proceedings of 17th European Conference on Artificial Intelligence, in conjunction with 4th Prestigious Applications of Intelligent Systems special section.

Castagnos, S., Brun, A. and Boyer, A. 2008. Probabilistic association rules for item-based recommender systems. In Proceedings of 4th European Starting AI Researcher Symposium, in conjunction with 18th European Conference on Artificial Intelligence.

Castagnos, S., Jones, N. and Pu, P. 2010. Eye-tracking product recommenders' usage. In Proceedings of 4th ACM Conference on Recommender Systems.

Castagnos, S., Brun, A. and Boyer, A. 2013. When diversity is needed... but not expected! In Proceedings of 3rd International Conference on Advances in Information Mining and Management.

Celma, O and Cano, P. 2008. From hits to niches? Or how popular artists can bias music recommendation and discovery. In Proceedings of ACM Conference on Recommender Systems.

Coskun, V., Ozdenizci, B., and Ok, K. 2015. The survey on near field communication. Sensors (Basel). 15(6): 13348–13405. http://www.ncbi.nlm.nih.gov/pmc/articles/PMC4507650/

Falk, J.H., Koran Jr., J.J., Dierking, L.D., & Dreblow, L. 2010. Predicting visitor behavior. Curator: The Museum Journal, 28: 249–258.

Fournier, R. and Viennet, E. 2015. Ammico : recommandation sociale pour la visite de musée. http://raphael.fournier-sniehotta.fr/files/AMMICO-AAFD-FournierViennet.pdf

Gavalas, D., and Kenteris, M. 2011. A web-based pervasive recommendation system for mobile tourist guides. Personal and Ubiquitous Computing, 15(7), pp. 759-770.

de Gemmis, M., Lops, P., Musto, C., Narducci, F., and Semeraro, G. 2015. Semantics-aware content-based recommender systems. In Recommender Systems Handbook, chapter 4.

Hidasi, B., Quadrana, M., Karatzoglou, A. and Tikk, D. 2016. Parallel recurrent neural network architectures for feature-rich session-based recommendations. In Proceedings of ACM International Conference on Recommender Systems.

Hijikata, Y., Shimizu, T. and Nishida, S. 2009. Discovery-oriented collaborative filtering for improving user satisfaction. In Proceedings of 13th International Conference on Intelligent User Interfaces.

Jannach, D., Zanker, M., Felfernig, A. and Friedrich, G. 2010. Recommender systems: An introduction. Cambridge University Press.

Koren Y., Bell R. M. and Volinsky C., 2009. Matrix factorization techniques for recommender systems, IEEE Computer, 42(8): 30-37.

L'Huillier, A., Castagnos, S. and Boyer, A. 2014. Understanding usages by modeling diversity over time. In Proceedings of 22nd Conference on User Modelling, Adaptation and Personalization.

L'Huillier, A., Castagnos, S. and Boyer, A. 2016. The new challenges when modeling context through diversity over time in recommender systems. In Proceedings of 24th ACM Conference on User Modelling, Adaptation and Personalization.

Liiv, I., Tammet, T., Ruotsalo, T., and Kuusik, A. 2009. Personalized context-aware recommendations in SMARTMUSEUM: Combining semantics with statistics. In Proceedings of 3rd International Conference on Advances in Semantic Processing, pp. 50-55.

Martín, J. and Trummer, C. 2005. SCALEX - A personalized multimedia information system for museums and exhibitions. In Proceedings of 1st International Conference on Automated Production of Cross Media Content for Multi-Channel Distribution (AXMEDIS), pp. 299-302.

Mathias, M., Moussa, A., Zhou, F., Torres-Moreno, J.M., Poli, M.S., Josselin, D., El-Beze, M., Carneiro Linhares, A. and Rigat, F. 2014. Optimisation using natural language processing: Personalized tour recommendation for museums. In Proceedings of Federated Conference on Computer Science and Information Systems, pp. 439-446.

McGinty, L., and Smyth, B. 2003. On the role of diversity in conversational recommender systems. In ICCBR. 276–290.

Mobasher, B., Jin, X., and Zhou, Y. 2004. Semantically enhanced collaborative filtering on the web. In Web mining: From Web to Semantic Web, pp. 57-76. Springer.

Naudet, Y., Grouès, V., Mignon, S., Arnould, G., Foulonneau, M., Djaghloul, Y., and Khadraoui, D. 2014. An approach for context-aware semantic recommendations in future Internet. In Semantic multimedia analysis and processing. Taylor & Francis.

Naudet, Y., Antoniou, A., Lykourentzou, I., Tobias, E., Rompa, J., Lepouras, G. 2015. Museum personalization based on gaming and cognitive styles: The BLUE experiment. International Journal of Virtual Communities and Social Networking, 7(2).

Noor, S. and Martinez, K. 2009. Using social data as context for making recommendations: An ontology based approach. In Proceedings of 1st Workshop on Context, Information and Ontologies.

Radlinski, F., and Dumais, S. 2006. Improving personalized web search using result diversification. In SIGIR '06: Proceedings of 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval.

Ricci, F., Rokach, F., and Shapira, B. 2015. Recommender Systems Handbook. Springer.

Ricci, F., Rokach, F., and Shapira, B. 2015. Recommender systems: Introduction and challenges. In Recommender Systems Handbook, Chapter 1.

Teevan, J., Dumais, S. T., and Horvitz, E. 2005. Personalizing search via automated analysis of interests and activities. In Proceedings of 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval.

Zhang, M., and Hurley, N. 2008. Avoiding monotony: improving the diversity of recommendation lists. In Proceedings of 2008 ACM Conference on Recommender Systems.

Zhou, S., Zhou, X., Yu, Z., Wang, K., Wang, H., and Ni, H. 2009. A recommendation framework towards personalized services in intelligent museum. In Proceedings of IEEE International Conference on Computational Science and Engineering, pp. 229-236.

Ziegler, C.-N., Lausen, G., and Schmidt-Thieme, L. 2004. Taxonomy-driven computation of product recommendations. In Proceedings of 2004 ACM CIKM International Conference on Information and Knowledge Management, pp. 406–415.

Ziegler, C.-N., McNee, S. M., Konstan, J. A. and Lausen, G. 2005. Improving recommendation lists through topic diversification. In Proceedings of ACM International Conference on the World Wide Web.

# 3. Machine learning, semantic reasoning and crowdsourcing

## 3.1.  Machine learning

### 3.1.1. Survey

*Machine learning* is a broad term that encompasses computer software that is able to learn —and thus improve its performance according to some measure— based on experience and without being explicitly programmed (Mitchell, 1997). Machine learning can be performed in a supervised fashion (using labelled training examples), in an unsupervised fashion (using unlabelled training data) or via reinforcement (using rewards from the environment embedded in the training data). A vast literature of machine learning approaches has been accumulated including reinforcement learning techniques (Sutton and Barto, 1998; Kaelbling et al., 1996), support-vector networks (Boser et al., 1992; Cortes and Vapnik, 1995), Bayesian networks (Pearl, 1988) and many more.

One of the most popular and successful artefacts for machine learning was inspired by neurobiology and the structure of the brain. An *artificial neural network* (ANN) is fundamentally "*a densely interconnected set of simple units*" (Mitchell, 1997). These units are named *neurons*, and can take multiple real-valued inputs and return a single real-valued output. In most cases, the output is the result of an activation function applied on the weighted sum of the inputs; popular activation functions include the threshold function, the sigmoid function or the linear function. An ANN is a set of neurons, and the network has a set of input neurons (where external variables are fed to the network) and a set of output neurons (which consist of deductions from the external variables). Popular examples of ANN use include pattern detection tasks (e.g. where the inputs are the pixel values of an image and the outputs can be the similarity of the image to a letter) and robot control (where the inputs may be the pixels of the road ahead of a vehicle and the outputs are a representation of the vehicle's steering direction) (Pomerleau, 1993). Despite the considerable success of many uses, the architecture of ANNs has been typically limited to a few layers (usually one or two) beyond the input and output layers. The preference for this "shallow" architecture was primarily due to the poor performance of deeper architectures with numerous layers of hidden neurons (Bengio et al., 2007). Recent work on deeper architectures has led to several breakthroughs in performance, however, by attempting to "*decompose the problem into subproblems and multiple levels of representation*" (Bengio, 2009). *Deep learning* can be achieved through energy-based models such as Restricted Boltzman Machines (Larochelle and Bengio, 2008; Smolensky, 1986) but also through back-propagation. The latter is achieved by training a deep network one layer at a time; essentially, the deep network is treated as a stack of shallow networks where the output of the previous shallow network is the input to the next (Vincent et al., 2010). Training such shallow networks sequentially is possible by treating each shallow network as an autoencoder (Hinton and Zemel, 1994).

Machine learning has often been applied to predict real-valued outputs of unseen data; preference learning is an instance of machine learning where the predictions are made regarding "*order relation on a collection of objects*" (Fürnkranz and Hüllermeier, 2010). Instead of comparing,

therefore, the set of real-valued outputs of an ANN using an object's parameters as input, preference learning directly outputs the order between such objects. A great strength of preference learning is that its training set, used for supervised learning, does not require explicit ratings of objects and can learn directly from their rankings. Such data is likely less noisy especially in tasks such as human self-reports in user survey questionnaires (Yannakakis and Hallam, 2011). Preference learning problems can be categorised as object ranking problems, where the goal is to predict the ranking of objects outside of a training set, and label ranking problems, where the goal is to find a general ranking of labels for all instances of a training set (Hüllermeier et al., 2008); each of these types of problems can be approached via modelling the relation directly or by indirect induction (Fürnkranz and Hüllermeier, 2010).

### 3.1.2. CrossCult applications

Since machine learning is a general tool for processing data, it can be applied to a multitude of tasks undertaken by the CrossCult project. Indicatively, supervised machine learning can be applied to any user modelling task, using the visitor's usage data (e.g. number of comments or locations visited in Pilot 4, or path followed in Pilot 1) as input and the visitor's performance as desired output, thus learning a mapping between interaction data and the visitor's reflection potential. Moreover, the new associations of users in e.g. pilot 4 can be used as desired output of an ANN which learns the mapping between the data (e.g. locations, threads) and the perceived association strength, creating new associations even in unseen data (i.e. data that have not been validated by user feedback). For instance, when a new POI is inserted by an expert, based on the tags, location and even text description a predictive model (ANN) can specify what the anticipated appeal (e.g. rating level or visit probability) for this POI is likely to be, thus aiding experts in designing new content and POIs.

Deep learning algorithms can also be used for similar purposes in CrossCult. For instance, deep learning for image recognition could be used for encoding images of paintings for Pilot 1 in an unsupervised manner, or in a supervised manner by using image groupings provided by users. However, due to the extensive volume of data needed for training such deep networks and the extensive computational effort of learning, deep learning will be considered for CrossCult if the data acquired warrants it.

Finally, preference learning is an alternative method of learning from data, and is ideal for user responses in the form of preferences. Thus, preference learning can be used in CrossCult for deriving user models of preference: similar to machine learning, the input can be gameplay performance (such as answers to questions, speed of movement etc.) and the output can be user rankings of e.g. points of interest in pilot 4. The learned model can then be used to predict how other users with similar gameplay will like these points of interest, leading to better recommendations and user needs anticipation inside the application.

### 3.1.3. Software, tools and algorithms

**Table 4: Machine learning software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
| --- | --- | --- | --- |

| Weka 3 | GNU GPL | Collection of machine learning algorithms for data mining. | As Weka algorithms can be called from arbitrary Java code, they can be integrated to any RESTful webservice for mining player usage data (e.g. for predictive tasks) |
|---|---|---|---|
| TensorFlow | Apache 2.0 open source license. | General computation library, used primarily for machine learning and deep neural networks research. | Depending on the amount of data, deep learning (for large amounts of data, e.g. knowledge bases) or traditional machine learning (for smaller datasets, e.g. player data) can be used in a distributed manner via TensorFlow. |
| Preference Learning Toolbox | GNU Lesser GPL | Software with graphic user interface for performing preference learning (machine learning that handles datasets with ordinal relations). | Preference learning is an ideal alternative to ratings and any comparisons of player data can be converted into rankings and used for classification, predictive tasks and recommendations. |
| Scikit-Learn | BSD License | Python library for machine learning (including regression, principal component analysis, classification, and much more). | The API and tutorials for this library which is still actively in development can be used directly for machine learning tasks as part of the back-end of CrossCult. |

## 3.1.4. References

Bengio, Y. 2009. Learning deep architectures for AI. Foundations and Trends in Machine Learning 2:1–127.

Bengio, Y., Lamblin, P., Popovici, D., and Larochelle, H. 2007. Greedy layer-wise training of deep networks. In Proceedings of Advances in Neural Information Processing Systems (NIPS), pp. 153–160.

Boser, B. E., Guyon, I. M., and Vapnik, V. N. 1992. A training algorithm for optimal margin classifiers. In Proceedings of 5th Annual Workshop on Computational Learning Theory, pp. 144–152.

Cortes, C. and Vapnik, V. 1995. Support-vector networks. Machine Learning 20:273–297.

Fürnkranz, J. and Hüllermeier, E. 2010. Preference learning, pp. 789–795. In Encyclopedia of Machine Learning. Springer US.

Haykin, S. 1998. Neural Networks: A Comprehensive Foundation. Prentice Hall PTR, 2nd edition.

Hinton, G. E. and Zemel, R. S. 1994. Autoencoders, minimum description length, and Helmholtz free energy. In Advances in Neural Information Processing Systems, pp. 3–10.

Hüllermeier, E., Fürnkranz, J., Cheng, W., and Brinker, K. 2008. Label ranking by learning pairwise preferences. Artificial Intelligence 172:1897–1916

Kaelbling, L. P., Littman, M. L., and Moore, A. W. 1996. Reinforcement learning: A survey. Journal of Artificial Intelligence Research 4:237–285.

Larochelle, H. and Bengio, Y. 2008. Classification using discriminative restricted Boltzmann machines. In Proceedings of 25th International Conference on Machine Learning, pp. 536–543.

Mitchell, T. 1997. Machine Learning. McGraw Hill.

Pearl, J. 1988. Probabilistic reasoning in intelligent systems. Morgan Kaufmann.

Pomerleau, D. 1993. Knowledge-based training of artificial neural networks for autonomous robot driving. In Robot Learning. Springer Science + Business Media, LLC.

Smolensky, P. 1986. Information processing in dynamical systems: Foundations of harmony theory, pp. 194–281. In Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1.

Sutton, R. S. and Barto, A. G. 1998. Introduction to Reinforcement Learning. MIT Press.

Vincent, P., Larochelle, H., Bengio, Y., and Manzagol, P.-A. 2008. Extracting and composing robust features with denoising autoencoders. In Proceedings of 25th International Conference on Machine learning, pp. 1096–1103.

Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., and Manzagol, P.-A. 2010. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. Journal of Machince Learning Research 11:3371–3408.

Yannakakis, G. and Hallam, J. 2011. Rating vs. preference: A comparative study of self-reporting, pp. 437–446. In Affective Computing and Intelligent Interaction, volume 6974, Lecture Notes in Computer Science. Springer.

## 3.2.    Semantic reasoning

### 3.2.1. Survey

Semantic reasoning is the process of deriving facts that are not explicitly expressed in an ontology or knowledge base. The tasks that semantic reasoners are able to perform include: (i) concept satisfiability, i.e. determining whether an individual can exist as an instance of a given class; (ii) concept subsumption, i.e. determining whether given two ontology classes A and B, A subsumes B; (iii) consistency checking, i.e. checking whether the descriptions of the individuals comply with the class and property definitions and the axioms of the ontology; (iv) instance checking, i.e. determining whether an individual is an instance of a given class; (v) query answering, i.e. returning the answers of a query with respect to a given ontology; (vi) realisation of an individual, i.e. finding all classes that an individual belongs to.

Most modern description logic reasoners use tableau algorithms to reason with OWL (or more generally Description Logic) ontologies. A tableau algorithm reduces all the above reasoning tasks to the concept satisfiability problem. For example, proving that A is subclass of B is reduced to proving that $A \sqcap \neg B$ is unsatisfiable, where $\sqcap$ stands for intersection and $\neg B$ denotes the negation (complementary class) of *B*. It then solves the concept satisfiability problem by attempting to construct an interpretation of the ontology in which the concept in question is satisfied (contains at least one individual); failing to find one such interpretation means that the concept is unsatisfiable. More details about how the tableau algorithms work can be found in (Baader and Sattler, 2001). Apart from the reasoning algorithms they use, other ways to characterise a reasoner are with respect to (i) other reasoning features such as soundness and completeness, expressivity and rule support; (ii) practical usability characteristics, e.g. supported platforms, availability and licence; and (iii) performance indicators, i.e. space and time complexity for certain reasoning tasks (Dentler et al., 2011).

Logical rules provide a more a more expressive form of reasoning over ontologies. Compared to ontology languages, which are designed to describe concepts of the application domain, rule languages are based on more expressive logics to enable modelling additional aspects of an application, such as policies, integrity constraints, exceptions and others. The integration of rules with ontologies can be achieved either by combining existing ontology and rule languages (heterogeneous integration) or by defining new ones, expressive enough to define rules, ontologies and their interactions (homogeneous integration). Prominent examples of the first category include: TRIPLE (Sintek and Decker, 2002), a rule language based on F-Logic that admits queries to the ontology in the bodies of rules; SWI Prolog, a logic programming system with a Semantic Web library which makes it possible to invoke RDF Schema and OWL reasoners from Prolog programs; and dl-programs (Eiter et al., 2008), which combine logic programming under the answer set semantics with Description Logics. Examples of homogeneous integration include: DLP (*Description Logic Programs*) (Grosof et al., 2003), a language obtained by the intersection of a Description Logic with Datalog rules; SWRL (*Semantic Web Rule Language*), which extends OWL with rules interpreted as First Order Logic implications; and DR-Prolog (Antoniou and Bikakis, 2007), which combines OWL ontologies and Defeasible Logic rules.

## 3.2.2. CrossCult applications

One of the reasons for choosing OWL as the language for the CrossCult ontology is that, being based on *Description Logics*, it enables semantic reasoning to be performed automatically. This will be needed at different stages of CrossCult. During the ontology development stage (Task 2.3), an OWL reasoner will be used to check the correctness of the CrossCult ontology, e.g. all ontology classes are satisfiable (they can admit at least one individual). While populating the ontology with individuals for the four pilots (Task 2.2), we will perform consistency checking to ensure that the descriptions of the individuals are valid with respect to the ontology definitions and axioms. Moreover, we will use the OWL reasoner as a mechanism for association discovery, i.e. to reveal "hidden" relationships between artefacts, cultural heritage items, artists or any other related entities associated to the same or different pilots, through entailment (this comes in addition to the knowledge discovery mechanisms discussed in Section 4.3). Using query answering, the reasoner will also perform any information retrieval tasks required by the pilot apps, as well as it will underpin recommendation processes. For example, depending on the way the user preferences are modelled (as relationships between the user and other ontology concepts or as rules), we will use an OWL reasoner or a rule-based reasoner to match a visitor's profile with specific parts of the pilots' collections or with other visitors with similar preferences. Additionally, the tools for experts to design pilot experiences will be equipped with semantic reasoning mechanisms to aid in the identification of relevant multimedia contents or concepts/topics to address in the games or visitor paths. References for the latter will be sought in previous publications by CrossCult partners such as (Blanco-Fernández et al., 2011).

## 3.2.3. Software, tools and algorithms

An extended survey of available Semantic Web reasoners is available in (Mishra and Kumar, 2001). From the great range of available solutions, those that fit better the project's needs and requirements are presented in Table 5.

**Table 5: Semantic reasoning software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|------|---------|-------------|--------------|
| **Hermit** | GNU GPL | Very efficient OWL reasoner, which can determine whether a given ontology is consistent and identify concept subsumption relationships, among other features. It is one of the default reasoners of Protégé, the ontology editor we use for the development of the CrossCult ontology. | We will use it during the ontology development stage to check the correctness of the CrossCult ontology. |
| **FaCT++** | GNU GPL | New generation of the OWL DL reasoner FaCT. It supports OWL DL and a subset of OWL 2 (the version of OWL we will use for the CrossCult ontology). FaCT++ is implemented in C++ and is very efficient due to its optimised tableau algorithms. It is one of the default reasoners of Protégé. | We will use it during the ontology population and the pilot development stages to perform the more complex reasoning tasks (e.g. entailment) on the CrossCult ontology. |
| **SWRLTab** | BSD 2-clause License | Protégé plugin that enables writing and applying SWRL rules to an OWL ontology. | We will possibly use it for representing the visitors' preferences and for matching them with the pilots' collections. |

## 3.2.4. References

Antoniou, G. and Bikakis, A. 2007. DR-Prolog: A system for defeasible reasoning with rules and ontologies on the Semantic Web. IEEE Transactions on Knowledge and Data Engineering 19(2): 233–245.

Baader, F. and Sattler, U. 2001. An overview of tableau algorithms for Description Logics. Studia Logica 69(1): 5-40.

Blanco-Fernández, Y., López-Nores, M., Pazos-Arias, J.J. and García-Duque, J. 2011. An improvement for semantics-based recommender systems grounded on attaching temporal information to ontologies and user profiles. Engineering Applications of Artificial Intelligence, 24(8):1385-1397.

Dentler, K., Cornet, R., ten Teije, A., and de Keizer, N. 2011. Comparison of reasoners for large ontologies in the OWL 2 EL profile. Semantic Web Journal 2(2): 71-87.

Eiter, T., Ianni, G., Lukasiewicz, T., Schindlauer, R., and Tompits, H. 2008. Combining answer set programming with description logics for the Semantic Web. Artificial Intelligence 172(12-13): 1495–1539.

Grosof, B., Horrocks, I., Volz, R., and Decker, S. 2003. Description Logic programs: Combining logic programs with Description Logic. In Proceedings of 12th International Conference on the World Wide Web: 48-57.

Mishra, R.B. and Kumar, S. 2011. Semantic Web reasoners and languages. Artificial Intelligence Review 35(4): 339-368.

Sintek, M. and Decker, S. 2002. TRIPLE - A query, inference, and transformation language for the Semantic Web. In Proceedings of 1st International Semantic Web Conference, pp. 364-378.

## 3.3.   Crowdsourcing

### 3.3.1. Survey

Crowdsourcing is defined as a participative, typically online, activity in which an individual, institution or company proposes to a group of individuals of varying knowledge, heterogeneity and numbers, via a flexible open call, the voluntary undertaking of a task (Estellés-Arolas and González-Ladrón-de-Guevara 2012). This undertaking always entails mutual benefit. Participating users receive the satisfaction of a certain need, which can be economic, social recognition, self-esteem or skill development. The crowdsourcer obtains and uses to their advantage what the users brought to the venture, and the form of the latter depends on the type of the activity undertaken. Crowdsourcing can be distinguished into explicit, which allows participants to create work for the advertised purpose, and implicit, which assumes that they work towards solving a problem as a side effect of another activity.

Two main types of crowdsourcing tasks can be distinguished: *micro-tasks* (Kittur, Chi, and Suh 2008; Musthag and Ganesan 2013) and *macro-tasks* (Haas et al. 2015). Micro-tasks are sets of small units of work that together comprise a large job. They usually require minimal levels of expertise, can be accomplished in very short time (seconds or minutes) and are performed in parallel by a large crowd. Examples include image recognition, sentiment analysis and data collection. Micro-task based crowdsourcing is suitable for jobs that can be easily decomposed to individual-level contributions, which can then be aggregated automatically to give an objectively assessable solution (e.g. an image can be correctly recognised or not). Macro-tasks, on the other hand, are larger-sized work units, which require expert skills, assume varying degrees of knowledge over a topic, may take more worker time and often involve task dependency, i.e. workers building on each other's contributions. Examples include writing a news article, designing a product logo or crafting a piece of software.

The application areas of crowdsourcing are many, and of varied purposes. Crowdsourcing can be used for *opinion seeking*, where the crowd is used for input and suggestions (with platforms such as SurveyMonkey), and *data collection*, where crowd inputs are used to improve a knowledge base (Ipeirotis and Gabrilovich 2014). A third area is *content creation*, for example designing an advertisement or video (Roth and Kimani 2014), using platforms like Tongal). The crowd work model can also be applied on *solution finding* where the crowd is asked to solve a complex problem (like climate change mitigation (Malone 2011) or ways to improve medicine practices) through platforms such as InnoCentive and BrightIdea. It can be used for design competitions where crowd participants submit designs (e.g. a product logo) for selection by the customer through platforms like 99designs or LogoTournament, and, more broadly, for innovation contests, where companies appeal to the crowd for ideas either directly (e.g. Lego Mindstorms

(Majchrzak and Malhotra 2013) or through innovation brokers (Simula and Ahola 2014). Specific application areas are also implemented and they include programming and software testing (with platforms such as TopCoder and uTest respectively), customer service, where knowledgeable workers provide help to entire customer bases (Alison 2011) (through platforms like Mila), and finally for raising campaign funds and micro-angel investing through *crowdfunding* platforms like Kickstarter.

In parallel to applications, research is exploring ways to improve the provision of crowdsourcing. Among the most notable methods we find workload and task assignment optimisation algorithms (Basu Roy et al. 2015; Bernstein et al. 2012; Ho and Vaughan 2012; Goel, Nikzad, and Singla 2014; Karger, Oh, and Shah 2014), the high-level goal of which is to provide performance guarantees to the crowdsourcing process. The idea is to allocate worker assignments per task such as the worker contributions add up to a required quality threshold within a given budget. Roughly speaking, the crowdsourced tasks play the roles of multiple knapsacks with some additional concepts like domain-specific expertise and wages per worker, different models of acceptance probabilities or types of quality aggregation. Incentive engineering approaches have also been proposed, with the aim of identifying on the one hand the appropriate motivations for crowdsourcing users and on the other hand the appropriate pricing mechanisms that will sustain and improve motivation to participate (Faridani, Hartmann, and Ipeirotis 2011; Minder et al. 2012; Zhang and Van Der Schaar 2012; Kaufmann, Schulze, and Veit 2011). Finally, interdisciplinary approaches drawing from computer science and phychology to improve crowd team collaboration and matchmaking have also been recently proposed (Yue, Ali, and Wang 2015; Lykourentzou et al. 2016).

### 3.3.2. CrossCult applications

In the context of CrossCult, crowdsourcing will be used as a means of: (i) acquiring user input in the form of ratings, text or other content (e.g. images) and (ii) processing user input, with the purpose of curating, augmenting and improving it, as well as to drive decision-making for other CrossCult technologies such as recommendations.

In regards to acquiring user input, we will apply both explicit crowdsourcing (e.g. visitors of the city using the CrossCult app rate POI descriptions and/or other visitor comments) and implicit (e.g. visitor paths are acquired to help better classification of items into categories). In regards to processing user inputs, we further distinguish two processing method families. First, for the processing of numerical input (e.g. ratings), crowdsourcing mechanisms can be used to assess the reliability of the rating providers and weigh their contribution, in order to reach an aggregated crowd opinion on a specific item (for example the description of a POI or museum item) or user content (for example a photograph taken by a visitor). Indicative algorithms that can be used towards this goal include majority voting, or more advanced methods of reliability estimation from noisy input and without the presence of a golden standard based on probabilistic machine learning (Raykar et al., 2010; Whitehill et al., 2009) or active learning (Yan et al., 2011). Second, for the processing of non-numerical user input (e.g. text or images), micro-task platforms (such as Amazon Mechanical Turk or CrowdFlower) can be used to process batches of anonymised user content and filter it (for example for hate speech, or inappropriate language), curate or

summarise it. Such crowdsourced models for recognising text could also potentially be used for recommendations and even user modelling (discussed in Section 2). Further, collaborative crowdsourcing techniques such as the ones proposed in (Lykourentzou et al., 2016a,b; Basu Roy et al., 2015) can also be used to summarise and augment the received content.

### 3.3.3. Software, tools and algorithms

**Table 6: Crowdsourcing software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|---|---|---|---|
| **Pybossa** | GNU Affero General Public License: free, copyleft license for software and other kinds of works, specifically designed to ensure cooperation with the community in the case of network server software. | PyBossa is an open source platform for crowdsourcing online (volunteer) assistance to perform tasks that require human cognition, knowledge or intelligence (e.g. image classification, transcription, information location etc). It can be used for any distributed tasks project but was initially developed to help scientists and researchers crowdsource human problem-solving skills! | Building blocks and front-ends to create a crowdsourcing portal. |
| **Hive** | Apache License, version 2.0 | An open-source platform that lets developers produce crowdsourcing applications for a variety of contexts. | Building blocks and front-ends to create a crowdsourcing portal. |
| **Ventour** | Propietary of GVAM | App builder and content management system for guiding apps in museums. It permits users to add textual content to their "my visit at" module | Comments gathering within guiding apps and retrieval system via email |
| **Vocalia** | SaaS | Vocalia is a natural languages recognizer that allows process complex text and permits interaction with voice control for apps. | Comments and crowdsourced text analysis to enhance recommendation and profiling. |
| **CrowdFlower** | SaaS | Micro-task platforms for repetitive short tasks. | Content filtering, sentiment analysis. |
| **Amazon Mechanical Turk** | SaaS | Micro-task platforms for repetitive short tasks. | Content filtering, sentiment analysis. |
| **Algorithms described in Raykar et al., 2010, Whitehill et al., 2009, Yan et al., 2011 (indicatively)** | N/A | Algorithms for learning rater reliability without the presence of a golden standard. | Estimate reliability of user ratings to improve decision-making regarding information classification in CrossCult platform/apps. |
| **Methods described in Lykourentzou et al.,** | N/A | Algorithms and methods for collaborative crowdsourcing. | Enable augmentation of textual user content by user- |

| 2016a,b, Basu Roy et al., 2015 (indicatively) | | | to-user or task-to-task matching. |
| --- | --- | --- | --- |

## 3.3.4. References

Alison, D. 2011. Oracle integrates crowdsourcing into CRM. InformationWeek. http://www.informationweek.com/applications/oracle-integrates-crowdsourcing-into-crm/d/d-id/1096652?

Basu Roy, S., Lykourentzou, I., Thirumuruganathan, S., Amer-Yahia, S. and Das, G. 2015. Task assignment optimization in knowledge-intensive crowdsourcing. VLDB Journal 24(4).

Bernstein, M.S., Karger, D.R., Miller, R.C. and Brandt, J. 2012. Analytic methods for optimizing realtime crowdsourcing. Arxiv Preprint arXiv: http://arxiv.org/abs/1204.2995.

Estellés-Arolas, E. and González-Ladrón-de-Guevara, F. 2012. Towards an integrated crowdsourcing definition. Journal of Information Science 38(2): 189–200.

Faridani, S., Hartmann, B. and Ipeirotis, P.G. 2011. What's the right price? Pricing tasks for finishing on time. Human Computation, 26–31.

Gagan, G., Nikzad, A. and Singla, A. 2014. Allocating tasks to workers with matching constraints: Truthful mechanisms for crowdsourcing markets. In Proceedings of 23rd International Conference on World Wide Web, 279–80.

Haas, D., Ansel, J., Gu, L. and Marcus, A. 2015. Argonaut: Macrotask crowdsourcing for complex data processing. In Proceedings of VLDB Endowment 8(12): 1642–53.

Chien-Ju, H. and Wortman Vaughan, J. 2012. Online task assignment in crowdsourcing markets. In Proceedings of 26th AAAI Conference on Artificial Intelligence, 45–51.

Ipeirotis, P. G. and Gabrilovich, E. 2014. Quizz: Targeted crowdsourcing with a billion (potential) users. In Proceedings of 23rd International Conference on World Wide Web: 143–54.

Karger, D. R., Oh, S. and Shah, D. 2014. Budget-optimal task allocation for reliable crowdsourcing systems. Operations Research 62(1): 1–24.

Kaufmann, N., Schulze, T. and Veit, D. 2011. More than fun and money. Worker motivation in crowdsourcing – A study on Mechanical Turk. In Proceedings of 27th Americas Conference on Information Systems, no. 2009: 1–11.

Kittur, A., Chi, E.H. and Suh, B. 2008. Crowdsourcing user studies with Mechanical Turk. In Proceedings of 2008 SIGCHI Conference on Human Factors in Computing Systems, 453–56.

Lykourentzou, I., Antoniou, A., Naudet, Y. and Dow, S.P. 2016a. Personality matters: Balancing for personality types leads to better outcomes for crowd teams. In Proceedings of ACM Conference on Computer Supported Cooperative Work (27).

Lykourentzou, I., Wang, S., Kraut, R.E. and Dow, S. P. 2016b. Team dating: A self-organized team formation strategy for collaborative crowdsourcing, In Proceedings of CHI Conference Extended Abstracts on Human Factors in Computing Systems, pp. 1243-1249.

Majchrzak, A., and Malhotra, A. 2013. Towards an information systems perspective and research agenda on crowdsourcing for innovation. Journal of Strategic Information Systems.

Malone, T. W. 2011. Solving climate change with crowdsourcing. MIT Sloan Experts Blog.

Minder, P., Seuken, S., Bernstein, A. and Zollinger, M. 2012. CrowdManager - Combinatorial allocation and pricing of crowdsourcing tasks with time constraints. Workshop on Social Computing and User Generated Content, in Conjunction with ACM Conference on Electronic Commerce.

Musthag, M. and Ganesan, D. 2013. Labor dynamics in a mobile micro-task market. In Proceedings of SIGCHI Conference on Human Factors in Computing Systems, pp. 641-650.

Raykar, V. C., Zu, S., Zhao, L. H., Valadez, G. H., Florin, C., Bogoni, L., Moy, L. 2010. Learning from crowds. Journal of Machine Learning Research, 11: 1297-1322.

Roth, Y. and Kimani, R. 2014. Crowdsourcing in the production of video advertising: The emerging roles of crowdsourcing platforms. In International Perspectives on Business Innovation and Disruption in the Creative Industries: Film, Video, Photography: 1–36.

Simula, H., and Ahola. T. 2014. A network perspective on idea and innovation crowdsourcing in industrial firms. Industrial Marketing Management 43 (3): 400–408.

Yan, Y., Rosales, R., Fung, G. and Dy, J.G. 2011. Active learning from crowds. In Proceedings of 28th International Conference on Machine Learning.

Tao, Y., Ali, S. and Wang, S. 2015. An evolutionary and automated virtual team making approach for crowdsourcing platforms. In Crowdsourcing: Cloud-Based Software Development: 113–30.

Whitehill, J., Wu, T., Bergsma, J., Movellan, J.R. and Ruvolo, P.L. 2009. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. Advances in Neural Information Processing Systems (22).

Yu, Z. and van der Schaar, M. 2012. Reputation-based incentive protocols in crowdsourcing applications. In Proceedings of IEEE INFOCOM: 2140–2148.

# 4. Context mining and knowledge discovery

## 4.1.    Context awareness and the Internet of Things

### 4.1.1. Survey

Context-aware systems are able to adapt their operations to the current context without explicit user intervention and thus aim at increasing usability and effectiveness (Baldauf, Dustdar and Rosenberg, 2007). It is challenging to define the term "*context*" and many researchers have provided their own definitions for what notions it actually includes. In literature, the term *context awareness* first appeared in (Schilit and Theimer, 1994), where the authors described context as location, identities of nearby people, objects and changes to those objects. Ryan et al. (1997) referred to context as the user's location, environment, identity and time. Dey (1998) defined context as the user's emotional state, focus of attention, location and orientation, date and time, as well as objects and people in his/her environment.

Nowadays many mobile systems react to changes of location, time and other environment attributes, retrieving context information in a variety of ways, such as applying sensors, network information, device status, browsing user profiles and using other sources (Liu, Li and Huang, 2011). While there the numerous sensors on mobile devices contextual information from sensors on user's mobile devices, location information is by far the most frequently used attribute of context, and at the time there are 5 mainstream technologies for location tracking available in the market: GPS, WiFi, beacons, RFID and RFC. The following is a quick overview of each of them with a rating for specific characteristics –retrieved from: http://lighthouse.io/indoor-location-technologies-compared/– , with 1 meaning the worst grade and 5 the best. The list includes TILT, the LIST proprietary Wi-Fi based location tracking solution. Future technologies to bear in mind include LED beacons[1], audio beacons (Javed, 2013) and LIFI systems[2].

| | Beacons | GPS | WiFi | TILT | RFID | NFC |
|---|---|---|---|---|---|---|
| **How it works** | Bluetooth low energy beacons send a signal; device detects signal and acts based on data service rules. | Satellite radio signals. GPS devices receive the signal and determine location. | Wireless access points detect devices and triangulate distance based on received signal strength. | Wi-Fi based tool using a fingerprinting method. | Radio 'tags' transmit stored information (passive or active) to 'readers' that record data and/or perform actions based on reader application | Passive UHF RFID chips (usually built into device or card) transmit data to terminals upon close contact. |

---

[1] http://www.ledsmagazine.com/articles/2016/03/osram-forays-into-indoor-positioning-starting-a-tech-battle.html (retrieved October 17th, 2016)

[2] http://www.directionsmag.com/entry/li-fi-offers-new-technology-for-indoor-positioning/386499 (retrieved October 17th, 2016)

| | | | | | software rules. | |
|---|---|---|---|---|---|---|
| **Typical range** | 1 - 50 meters | Unlimited | 20 - 50 meters | 20 - 50 meters | 1cm - 100m | 10cm or less |
| **Network infrastructure requirements** | 4 | 4 | 3 | 3 | 2 | 3 |
| **Accuracy** | 4 | 1 | 3 | 4 | 5 | 5 (near range only) |
| **Privacy & security** | 3 | 3 | 2 | 3 | 4 | 4.5 |
| **Cost** | 4 | 5 | 3 | 3 | 2 | 3 |
| **Best for** | Indoor tracking; passive notification of contextual information; peer-to-peer messaging. | Outdoor tracking and navigation; agriculture and military uses. | Existing infrastructure and/ or strong need for WiFi connection and location information accuracy is only required within meters. | Existing infrastructure and/ or strong need for WiFi connection and location information accuracy is required within meters. | SKU level tracking of inventory, requirements for centimetre accuracy. | One-to-one secure delivery of information between consumer and another entity (payment, ticketing, etc.) |

For the purpose of the project, we can exclude RFID as it requires specific hardware at the receptor's/visitors mobile devices and infrastructure to work. NFC is not a good fit either because of its (very) short range nature.

GPS access is provided natively by all modern mobile OS and most smartphone feature GPS hardware; however, more and more mobiles at medium and low band of price do not have a proper GPS sensor. Mobile operating systems provide out of the box location services that don't rely on GPS data but can take advantage of cellular network information and Wi-Fi hotspots location to obtain a more accurate location. The exact use of the different sources is platform specific and this kind of solution is specifically targeting outdoor scenarios.

In indoor scenarios where the use of GPS is not possible, the first source of information regarding position is Wi-Fi. Indeed, Wi-Fi hotspots tend to be available in most public places and venues. A common technique used to create a Wi-Fi positioning system is fingerprinting, i.e. measuring the intensity of the received signal from the user's device and compare the data with a position database (Duque Domingo et al., 2016), then sending the information back to the device to "paint" it on a map or inform the user. This is the approach used in TILT.

The emergence of Bluetooth Low Energy technology (BLE) allowed the development of BLE beacons, small hardware transmitters that broadcast messages in a certain range. Beacons are meant for proximity-based interactions. Knowing the location of a beacon, it is possible to infer the location of a device based on the strength of the received signal. Then, using multiple beacons, triangulation/trilateration techniques can provide more accurate locations. The advantages of beacon technology are power efficiency and low cost (a beacon costs only few dollars). Currently, there exist two main product lines: iBeacon (Apple's proprietary specification, compatible with Android 4.3 and later) and Eddystone (Google-supported and open-source). An extensive comparison of the two technologies can be found in (Mittal, 2016).

Modern studies on context awareness are inspired by the concept of the Internet of Things (IoT) —"*thing*" in this case is an "*inextricable mixture of hardware, software, data and service*" (Noto La Diega

and Walden, 2016). The main goal of IoT is to allow objects to be sensed and/or controlled remotely across existing network infrastructure, which finds applications ranging from everyday life in smart buildings and smart cities, to the industry of the future. In the IoT vision, the systems combine physical machinery with a cyber presence (Broy, 2011), and can follow the 5C (connection, conversion, cyber, cognition, configuration) architecture of (Lee, Bagheri and Kao, 2015). This hierarchical design architecture includes different levels of technical implementation, with higher levels requiring (and enriching) elements in the lower levels. The 5C levels include (i) the *Smart Connection Level*, acquiring data from devices, (ii) the *Data-to-Information Conversion Level*, converting raw data into meaningful, human-interpretable results such as health assessment, (iii) the *Cyber Level*, linking and comparing against data from other devices or historical trends, (iv) the *Cognition Level*, of devices being aware of their potential failures, and (v) the *Configuration Level*, with machines adjusting their working load or schedule to address discovered potential failures. Data acquired from many device sensors (GPS, accelerometers, gyroscopes, magnetometers, proximity sensors, light sensors, even barometers, thermometers, heart rate monitors, etc.), in conjunction with any other features of context fed from whichever source, make up an extremely varied aggregate that is being dealt with in *sensor fusion* approaches (Martí Muñoz, 2015).

## 4.1.2. CrossCult applications

Two of the CrossCult pilots (pilot 1 and pilot 4) rely extensively on context identification as a driver of the experiences. Pilot 1 is about indoor-location tracking, whereas Pilot 4 happens outdoors, so they induce different constraints (e.g. GPS cannot be used indoors). Ideally, the solution developed as part of the CrossCult framework should work seamlessly in indoor & outdoor environments.

For the task of context mining, CrossCult will focus on the first and second levels of the 5C model. Mobile devices have numerous sensors that provide valuable contextual information. Going beyond the 'pure' output of a GPS sensor, its traces can be analysed over time in order to derive their speed of movement; this has been used to motivate for instance exercising for games which target physical activity, such as Pokemon Go (2016), where if a device's GPS signal changes location drastically within minutes would indicate that the user is not walking but instead using other means of transportation. Similarly, detecting a lack of movement from the GPS signals can indicate that the user has lost interest in the current task or is suffering from congestion; a micro-augmentation, a new task or a path recommendation that takes the user away from the congested area can counter this behaviour.

## 4.1.3. Software, tools and algorithms

**Table 7: Location tracking software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|---|---|---|---|
| **Open GPS Tracking System** | Apache License V2.0 | Project designed specifically to provide web-based GPS tracking services for a fleet of vehicles. | The protocols of this software can be used as inspiration or implementation for pilots which include large volumes of users moving simultaneously in a space. |
| **MapNav -** | Proprietary | MapNav is a Geolocation | For the case of unity-based game development, |

| | | | |
|---|---|---|---|
| **Geolocation Toolkit** | | Engine for Unity and a powerful tool to develop location-based mobile apps and games, including 2d/3d object geolocation, GPS navigation and online maps. | MapNav can be used for GPS navigation. |
| **HTML5 Geolocation API** | Free | The HTML Geolocation API is used to get the geographical position of a user. | In lightweight applications in HTML, the geolocation API is an alternative to bulkier systems above, provided that the website is in a secure context (HTTPS). |
| **Ventour** | Proprietary of GVAM | App builder and content management system for guiding apps in museums. Its positioning module control maps interfaces with related POIs (tag value for indoor positioning system or geo-coordinates for GPS systems) | Management of maps, POIs and trajectories for guiding apps |
| **Cordova geolocation plugin** | Apache 2.0 | Cordova plugin (cross-platform mobile development framework) to get information about a device's location. Uses any available location information sources (exact implementation is platform specific). | Straightforward use in outdoor scenario to obtain GPS coordinates of the user. |
| **Cordova geofence plugin** | Apache 2.0 | Cordova plugin (cross-platform mobile development framework) to create and manage geofences (i.e. physical areas that will trigger a notification when entered/exited) | Can be used to trigger content delivery in an outdoor scenario, event when the app is not launched. |
| **TILT** | LIST proprietary | A Wi-Fi based location system that allows to obtain GPS coordinates inside a venue. The system provides an Android SDK, location tracking back end features and leaflet integration for displaying the position in a map based UI. | Use as an indoor location information source for pilot 1. Can be mixed with other sources for better accuracy. |
| **Eddystone** | Apache 2.0 | An open Bluetooth 4.0 protocol from Google, with official support for both iOS and Android. | Use as an indoor location information source for pilot 1. Can be mixed with other sources for better accuracy. |
| **Estimote beacons** | Commercial | Both hardware and software solution to install and manage beacons in a venue. Compatible with both iBeacon & Eddystone. | Install in NG on a per room basis to detect the user entrance in a specific room. |

## 4.1.4. References

Baldauf, M., Dustdar, S., and Rosenberg, F. 2007. A survey on context-aware systems. International Journal of Ad Hoc and Ubiquitous Computing, 2(4): 263-277.

Broy. M. 2011. Cyber-physical systems: Innovation durch softwareintensive eingebettete Systeme. Springer-Verlag.

Dey, A.K. 1998. Context-aware computing: The CyberDesk project. In Proceedings of AAAI, Spring Symposium on Intelligent Environments, pp. 51–54.

Duque Domingo, J., Cerrada, C., Valero, E., and Cerrada, J.A. 2016. Indoor positioning system using depth maps and wireless networks. Journal of Sensors, 14:1-8.

Javed, T. 2013. Enabling indoor location-based services using ultrasound. Ph.D. Theses. Queen's University. Ontario, Canada.

Lee, J., Bagheri, B. and Kao, H.A. 2015. A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. Manufacturing Letters, 3: 18-23.

Liao, L., Patterson, D.J., Fox, D. and Kautz, H. 2005. Building personal maps from GPS data. In Proceedings of IJCAI Workshop on Modeling Others from Observation.

Liu, W., Li, X., and Huang, D. 2011. A survey on context awareness. In Proceedings of International Conference on Computer Science and Service System.

Martí Muñoz, E.D. 2015. A framework for context-aware sensor fusion. Ph.D. Thesis. Universidad Carlos III de Madrid.

Mittal, S. 2016. iBeacon vs Eddystone: Which one works better for your Pilot Project? http://blog.beaconstac.com/2016/01/ibeacon-vs-eddystone/

Noto La Diega, G. and Walden, I. 2016. Contracting for the Internet of Things: Looking into the Nest. Queen Mary School of Law Legal Studies Research Paper No. 219/2016.

Quan, M., Navarro, E., and Peuker, B. 2010. Wi-fi localization using RSSI fingerprinting.

Ryan, N., Pascoe, J. and Morse, D. 1997. Enhanced reality fieldwork: the context-aware archaeological assistant. In Proceedings of 25th Anniversary Computer Applications in Archaeology.

Schilit, B. and Theimer, M. 1994. Disseminating active map information to mobile hosts. IEEE Network, 8(5): 22–32.

Williams, J. 2014. Internet of Things: Science fiction or business fact (Report and Resources). Harvard Business Review.

## 4.2.  Open Data

## 4.2.1. Survey

The web is a vast repository of information that can provide contextual information between different datasets, or real-world knowledge about current events or trending topics of discussion. The Open Data initiative, in particular, allows for the coupling of datasets without restrictions and (more importantly) using standardised protocols for communication and data mining. In short, "*open data and content can be freely used, modified, and shared by anyone for any purpose*" (The Open Definition), and its benefits for contextual mining lie in the requirement that open data *"[…] must be provided in a form readily processable by a computer and where the individual elements of the work can be*

*easily accessed and modified*' (The Open Definition). Contextual mining can be used on open data in order to discover associations between dissimilar sources such as articles from Wikipedia, images from Wikimedia commons and maps from OpenStreet Maps (Barros, Liapis, Togelius, 2016). The Semantic Web has prioritised structured, in common data formats and languages specifically designed for describing data, such as OWL as surveyed in Section 3.2. Data originating from the same dataset and sharing the same format can be mined via more sophisticated approaches: examples include topic models which can learn the preference of a specific mobile device user from their historical context logs (Bao et al., 2010), finding geographic (e.g. proximity to restaurants) and temporal (e.g. time of day) in the GPS history of mobile device users (Liao, Patterson, Fox, Kautz, 2010), or finding common patterns in the visual depictions of avatars created by multiple players (Lim, Liapis, Harrel, 2016).

## 4.2.2. CrossCult applications

Open data repositories can be used for discovering contextual information such as the time of day, the weather, special events (or holidays) associated with this day and the location of the user. This contextual information can be used to adapt the experience of the user, to enrich the recommendations provided to the user, or to inform the user model. There are instances where games such as World of Warcraft (Blizzard 2004) have special events associated with major holidays such as Christmas and Easter, complete with new quests or rewards; this will be considered for the CrossCult pilots provided that the development effort does not come at the cost of other pilot requirements. The user model, on the other hand, can include information such as the time of day, weather, and other information: this can be used to personalize the recommendations to e.g. avoid open areas when suggesting paths (to avoid the rain) or prioritize visiting a location where there is a religious festival (in the case of a holiday).

Another feature of context awareness can be found in pilot 2, namely in the design of experiences by choosing the nodes to display and the multiple choice questions to include in the graphs of connections. In this case, an ontology of relevant international dates (e.g. the International Day of Water, the birthday of Maria Skłodowska-Curie, or the National Day of Hungary) will be created that will enable the machine learning mechanisms of the CrossCult platform to identify opportunities/motivations to explore links among different venues.

## 4.2.3. Software, tools and algorithms

**Table 8: Open data software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|------|---------|-------------|--------------|
| **Open Street Maps** | Open Data Commons Open Database License | Collection of data about roads, trails, cafés, railway stations, etc. all over the world. | The city data can be used to automatically detect (and possibly recommend) paths and/or items of interest for city-based applications (such as pilot 4). |
| **Wikimedia Commons** | Almost all content can be freely reused | Database of over 33 million media files (images, sounds, videos, etc.) | While the content of CrossCult applications will not be automatically scraped from sites such as Wikimedia, this is a great source for early searches of transmedia associations. |
| **WikiMystery** | N/A | Adventure game generated entirely from Wikipedia articles, Open Street maps | While the games developed under CrossCult will not be automatically generated, the system description can be used as inspiration for |

| | | and Wikimedia assets. | enriching human-designed games with some data-driven elements. |
|---|---|---|---|
| **RelFinder** | GNU General Public License | Allows users or software to It extract, visualize and explore relationships in RDF data, e.g. Wikipedia articles. | The RelFinder algorithm can be used as a basis for discovering contextual links between elements inserted by CrossCult users (venue experts or end-users) to enrich the user input. |

## 4.2.4. References

Bao, T., Cao, H., Chen, E., Tian, J. and Xiong, H. 2010. An unsupervised approach to modelling personalized contexts of mobile users. In Proceedings of IEEE International Conference on Data Mining, pp. 38–47.

Barros, G.A.B., Liapis, A. and Togelius, J. 2016. Playing with data: Procedural generation of adventures from Open Data. In Proceedings of the International Joint Conference of DiGRA and FDG.

Liao, L., Patterson, D.J., Fox, D., and Kautz, H. 2005. Building personal maps from GPS data. In Proceedings of IJCAI Workshop on Modeling Others from Observation.

Lim, C.U., Liapis, A., and Fox Harrell, D. 2016. Discovering social and aesthetic categories of avatars: A bottom-up artificial intelligence approach using image clustering. In Proceedings of the International Joint Conference of DiGRA and FDG.

The Open Definition. http://opendefinition.org/ Retrieved 4 September 2016.

Blizzard 2004. World of Warcraft Massive-Multiplayer Online Game.

## 4.3.    Knowledge discovery

## 4.3.1. Survey

The *Knowledge Discovery in Databases* process (KDD) consists in processing large volumes of data for discovering significant and reusable patterns. From an operational point of view, KDD is based on three main operations: preparation of the data, data mining, and finally interpretation of the discovered patterns. Moreover, the KDD process is iterative, interactive, and controlled by an analyst, possibly using domain knowledge and experience for improving the KDD process. There are many connections between knowledge discovery and knowledge processing, knowledge guiding the discovery process while the output of KDD may feed knowledge bases.

*Formal Concept Analysis* (FCA) is an efficient candidate method for KDD, which is based on classification and allows multiple tasks to be performed, e.g. for discovering, representing, and problem solving (Ganter and Wille, 1999; Poelmans et al., 2010). FCA starts with a formal context including a set G of objects, a set M of attributes, or items, and a binary relation I between G and M. Two derivation operators are acting dually, the first associating to a set of objects the set of attributes that they share, and the second to a set of attributes the set of objects to which they belong. The composition of two derivations defines a closure operator. Then a concept is a pair (A,B) where the extent A is a closed set of objects while the intent B is the closed set of corresponding attributes. The form of a concept as a pair is important as it is in

accordance with concepts in ontologies. The set of concepts can then be ordered within a complete lattice called the concept lattice thanks to the inclusion of the extents (or the dual inclusion of intents). Moreover, implications can be also extracted from the lattices (rules with a full confidence). The concept lattice is a very useful structure for visualisation, navigation, and information retrieval, provided that the original dataset is not too huge (Valtchev et al. 2004; Alam et al., 2016).

However, natural data are not always binary and more complex data should be also processed. Then the formalism of pattern structures generalizes the classical settings of FCA and can be applied to various kinds of data such as numbers, sequences, trees and graphs. Actually, the set M of attributes in plain FCA is replaced with a lattice of (complex) descriptions where there are two main operations, namely the similarity between two descriptions and an associated partial ordering which support the organisation of descriptions within a concept lattice.

Finally, *Relational Concept Analysis* (RCA) is another extension of FCA that allows to directly take into account relations between objects themselves (Rouane Hacene et al., 2013). Then concepts are composed of relational attributes that are connecting objects to concepts and concepts to concepts. Relational attributes can be considered as roles in Description Logics (DL), making RCA a good formalism for guiding ontology design (Guédi et al., 2013; Charnay et al., 2015).

Efficient algorithms exist for building concept lattices, pattern structures and relational structures. Thus FCA and its extensions are used in numerous application domains for various tasks. Among these tasks, information retrieval and recommendation are two tasks where FCA shows many capabilities (Codocedo et al., 2014; Codocedo and Napoli, 2015). In particular, FCA can be used for biclustering, i.e. finding concepts with complex regularities inside (Codocedo and Napoli, 2015). In this way, it can be adapted to collaborative filtering and recommendation, when the similarity is based on the closeness of two columns in a table, i.e. searching for all individuals rating an item in the same or similar way. Finally, another task in which FCA and pattern structures can be reused with very good possibilities is the discovery and the management of sequences (Buzmakov et al., 2016). There, two tasks such as the discovery of particular sequences and then the organisation of these sequences can be combined within one process.

## 4.3.2. CrossCult applications

Primarily, knowledge discovery methods will be used when parsing the databases (including cultural historical databases, venue agendas or even city agendas and local news) for selecting content pertinent to the CrossCult goals of re-interpreting European (hi)stories. The techniques described above have shown their potential not only in searching within such databases, which can be valuable for experts that prepare CrossCult content, but also in automatically discovering patterns among data entries. Discovered pairs of data entries can e.g. be used for Pilot 4, if one of the entries refers to Luxemburg and the other refers to Valletta. On the other hand, recommendations derived from automatically discovered regularities, patterns, or sequences of data can be provided either to venue experts or even end-users, thus allowing for the crowdsourcing of knowledge discovery. However, the potential of presenting automatically

discovered relationships within databases to end-users will greatly depend on the quality and quantity of such relationships, as well as the development effort required for creating an interface for evaluating, rating, or censoring such relationships by end-users. To be precise, if the relationships found via KDD methods are extremely vast and/or not representing intuitive associations that human users would understand, using the crowd would not be desirable as they would be constantly bombarded by large amounts of nonsensical relationships. On a more controllable level, discovered relationships could be verified, censored and sanitised by venue experts who have more expertise with the domain itself and such computational tools in general.

### 4.3.3. Software, tools and algorithms

**Table 9: Knowledge discovery software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|------|---------|-------------|--------------|
| **Conexp** | Freely available | The Conexp system is one of the most used for building and drawing lattices from a binary table for FCA purposes. The system can display lattices of medium size (with 1000 nodes at most) and integrates basic functionalities for navigation and search. | The Conexp system can be used within CrossCult for searching the datasets during development of the applications. |
| **Coron** | Registered at the "Agence pour la Protection des Programmes" (APP) and is freely available | The Coron system includes a complete collection of data mining algorithms for extracting patterns such as frequent and closed itemsets, association rules and implications, and rules bases. In addition, the system proposes modules for cleaning and reducing the input dataset, and for visualising the resulting patterns. | The Coron system can be used within CrossCult for cleaning and reducing the datasets during development of the applications. |
| **Galicia** | Freely available | The Galicia system is used for building and drawing lattices from a binary table for FCA purposes. The system can display lattices of medium size (up to 1000 nodes) and integrates various functionalities for visualisation, navigation and search. One of the main functionalities of Galicia is to be able to display relational concept lattices within the Relational Concept Analysis framework. | The Galicia system can be used within CrossCult for finding relational concept lattices within the datasets used during development of the applications. |
| **LatViz** | Freely available | The LatViz system is still under development and is used for building and drawing lattices for FCA purposes. The system can start from binary or numerical tables as it is based on FCAPS which is able to work with pattern structures. The system can display lattices of large size (until 5000 nodes) as it integrates functionalities for local and interactive visualisation (focus, iceberg techniques), navigation, search and annotation. | Once completed, the LatViz system can used for visualisation, navigation and search of the datasets during development of the CrossCult applications. |

### 4.3.4. References

Alam, M., Nguyen Le, T.N., and Napoli, A. 2016. Latviz: A new practical tool for performing interactive exploration over concept lattices. In Proceedings of 13th International Conference on Concept Lattices and their Applications, pp. 9–20.

Buzmakov, A., Egho, E., Jay, N., Kuznetsov, S.O., Napoli, A., and Raïssi, C. 2016. On mining complex sequential data by means of FCA and pattern structures. International Journal of General Systems, 45(2): 135–159.

Charnay, C., Lachiche, N. and Braud, A. 2015. Approche relationnelle de l'apprentissage de séquences. In Proceedings of 15èmes Journées Francophones Extraction et Gestion des Connaissance, pp. 481-482.

Codocedo, V., Lykourentzou, I., and Napoli, A. 2014. A semantic approach to concept lattice-based information retrieval. Annals of Mathematics and Artificial Intelligence, 72:169-195.

Codocedo, V. and Napoli, A. 2014. Lattice-based biclustering using Partition Pattern Structures. In Proceedings of 21st European Conference on Artificial Intelligence,.

Codocedo, V. and Napoli, A. 2015. Formal concept analysis and information retrieval – A survey. In Proceedings of 13th International Conference on Formal Concept Analysis, pp. 61–77.

Guédi, A.O., Miralles, A., Huchard, M. and Nebut, C. 2013. A practical application of relational concept analysis to class model factorization: Lessons learned from a thematic information system. In Proceedings of 10th International Conference on Concept Lattices and their Applications, pp. 9-20.

Ganter, B. and Wille, R. 1999. Formal Concept Analysis. Springer, Berlin.

Poelmans, J., Elzinga, P., Viaene, S. and Dedene, G. 2010. Formal concept analysis in knowledge discovery: A survey. In Proceedings of 18th International Conference on Conceptual Structures: From information to intelligence, pp. 139-153.

Rouane Hacene, M., Huchard, M. and Valtchev, P. 2013. Relational concept analysis: Mining concept lattices from multi-relational data. Annals of Mathematics and Artificial Intelligence, 67(1):81-108.

Valtchev, P., Missaoui, R., Godin, R. 2004. Formal concept analysis for knowledge discovery and data mining: The new challenges. Lecture Notes in Computer Science, 2961:352-371.

# 5. Sporadic social networks and crowd management

## 5.1.    Sporadic social networks

### 5.1.1. Survey

Research in information services has made significant progress in exploiting the knowledge contained (explicitly or implicitly) on social networks like Facebook, Twitter, Instagram or LinkedIn. Despite their largely different approaches and objectives, these Web 2.0 sites are based on semi-permanent relationships among people. These relations serve to gradually build knowledge bases in the form of graphs, whose analysis enables features like recommending potentially interesting contents for each individual, launching advertising campaigns aimed at specific groups or segments of the population, identifying affinities among people or synergies between different areas of activity, etc. However, the interactions enabled by those sites are largely confined to the virtual world of the Internet. They are not accompanied by actual, face-to-face interactions except in cases in which people communicate to arrange physical meetings for entertainment or work. Moreover, it is noticeable that the individuals' interactions are increasingly focused on the set of people included in their social graphs, which are now accessible at any time. This causes a side effect of *de-socialisation*, in which the individual is isolated from his/her environment and voluntarily (though perhaps not quite consciously) gets trapped in a bubble of communication with his/her contacts (Kuss and Griffiths, 2011). Notwithstanding, the current state of technology has led many authors to envisage a new era of information services tailored to the people's physical and social context, aimed at facilitating the creation and exploitation of *sporadic (short-lived) social networks* (SSNs) where each individual can communicate with the people that surround him/her at a given moment (both acquaintances and strangers) to exchange information that may be relevant to them in different contexts and at different levels (room, building, street, city, province, etc.) (Ben Nejma et al., 2013; Bravo-Torres et al., 2014; Zhou et al., 2012). The proposal is applicable in various areas, from the formation of groups and the orchestration of activities around events or venues that attract people with potentially-related interests (e.g. museums, concert halls or campsites) (Ben Nejma et al., 2015, Srba et al., 2015) to opportunities for enhanced communications and access to relevant information on the road (vehicular social networks) (Smaldone et al., 2008; Bravo-Torres et al., 2013) or advances in the vision of the smart cities, related to the planning of personal mobility or the celebration of location-based urban games (Ryan-Collins, Stephens, and Coote, 2013).

### 5.1.2. CrossCult applications

In the context of CrossCult, the concept of sporadic social network is the creation of virtual groups or physical gathering of visitors on-the-fly in cultural institutions. Visitor groups thus formed may have similar preferences or objectives, or otherwise demonstrate complementarity in terms of knowledge and expectations. Visitors of similar groups allow their members to discuss, share their experiences and their views in relation to an exhibit, a section of the museum or the museum experience as a whole. This shared experience can be used for recommendation purposes in real time, as discussed in Section 2.3. On the contrary, visitors having different

preferences, knowledge and behaviours can be assembled to facilitate the resolution of tasks (scenarios imagined in the context of serious games), or to encourage debate and reinterpretation of history around the themes defined by historians. To achieve these objectives, it is necessary to collect all available information concerning such visitors. This information allows the creation of a profile for each visitor representing preferences and characteristics (user profiling described in detail in section 2.1). Specifically, the estimation of user preferences relatively to the exhibits will be based on implicit criteria (consultation time, frequency, or how recent the consultation was) via *use mining* techniques (Castagnos, 2008). User modelling will fit on it in a holistic approach by integrating human factors inherent in decision-making (tolerance in the distance, the crowd, technological intrusion) (Osche et al., 2016). These factors can be treated as independent variables, or correlated ones. We will rest on probabilistic graphical models (Markov chains, Bayesian networks, etc.) (Castagnos et al., 2005), time models (L'Huillier et al., 2016), clustering techniques (Castagnos et al., 2007) or multi-linear regression (Castagnos, 2008) to estimate the internal state of each visitor at a time *t* and predict his/her future actions. Depending on the profile of each visitor, it is then possible to determine via collaborative filtering algorithms individuals being the most similar to each other in terms of route (Li et al., 2008). Similarly, it is possible to form heterogeneous groups of individuals on the basis of complementary dimensions mining models (Despotakis et al., 2013; Loh et al., 2010; Zhang and Bockstedt, 2016).

## 5.1.3. Software, tools and algorithms

**Table 10: Sporadic network software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|---|---|---|---|
| **SPORANGIUM** | Proprietary of UVIGO | A software platform that provides constructs for sporadic social networks in different scenarios (smart buildings, smart cities, vehicular networks), with solutions from low-level communications to resource sharing and service-building blocks. Under development. | Algorithms for profile matching and context-aware group formation. |
| **Ventour APPs** | Propietary of GVAM | Methods to share content in museum apps and track user interaction level in his social profiles. | Direct use in Content sharing in SSNN |
| **Popcorm** http://www.pewe.sk/try-it/#popcorm | Proprietary of Slovak University of Technology in Bratislava | PopCorm (Popular Collaborative Platform) is a collaborative environment developed to support effective collaborative learning by creating dynamic short-term groups, using a method based on students' various personal and collaborative characteristics. | Reference algorithms for profile matching and group formation. |
| **Kalimucho** | Proprietary; middleware licensing. | Kalimucho is a Java framework that allows adding, updating or removing core application features in real time, while the app is running and without any | Reuse features demonstrated in applications of ephemeral communities and thematic interactions. |

| | user action. | |
|---|---|---|

## 5.1.4. References

Ben Nejma, G., Roose, P., Gensel, J., Dalmau, M., and Amine Ghorbali, M. 2013. Taldea: A tool for fostering spontaneous communities. In Workshops of 13th International Conference on Web Engineering, pp. 293-301.

Ben Nejma, G., Roose, P., Dalmau, M., and Gensel, J. 2015. Service discovery for spontaneous communities in pervasive environments. In Proceedings of 16th International Conference on Web Information Systems Engineering, pp. 337-347.

Bravo-Torres, J.F., López-Nores, M., Blanco-Fernández, Y., and Pazos-Arias, J.J. 2013. Leveraging short-lived social networks in vehicular environments. In Proceedings of 2nd International Conference on Future Generation Communication Technologies, pp. 196-200.

Bravo-Torres, J.F., López-Nores, M., Blanco-Fernández, Y., and Pazos-Arias, J.J. 2014. A platform to exploit short-lived relationships among mobile users: A case of collective immersive learning. In Proceedings of 20th International Conference on Information and Software Technologies, pp. 384-395.

Castagnos, S., Boyer, A., and Charpillet, F. 2005. Vers un filtrage collaboratif distribué: le modèle RSB. 3èmes Journées Francophones sur les Modèles Formels de l'Interaction.

Castagnos, S. and Boyer, A. 2007. Personalized communities in a distributed recommender system. In Proceedings of 29th European Conference on Information Retrieval.

Castagnos, S. 2008. Modélisation de comportements et apprentissage stochastique non supervisé de stratégies d'interactions sociales au sein de systèmes temps réel de recherche et d'accès à l'information. Thèse de doctorat en Informatique, Université Nancy 2.

Despotakis, D., Dimitrova, V., Lau, L., and Thakker, D. 2013. Semantic aggregation and zooming of user viewpoints in social media content. In Proceedings of 21st International Conference on User Modeling, Adaptation and Personalisation.

Kuss, D., and Griffiths, M. 2011. Online social networking and addiction: A review of the psychological literature. International Journal of Environmental Research and Public Health, 8: 3528-3552.

L'Huillier, A., Castagnos, S., Boyer, A. 2016. Modéliser la diversité au cours du temps pour détecter le contexte dans un service de musique en ligne. Revue des Sciences et Technologies de l'Information - Série TSI : Technique et Science Informatiques.

Li, Q., Zheng, Y., Xie, X., Chen, Y., Liu, W., & Ma, W. Y. (2008, November). Mining user similarity based on location history. In Proceedings of the 16th ACM SIGSPATIAL international conference on Advances in geographic information systems (p. 34). ACM.

Loh, S., Lichtnow, D., Kampff, A., Cerveira, J., and Moreira de Oliveira, J.P. 2010. Recommendation of complementary material during chat discussions. Knowledge Management & E-Learning: An International Journal (KM&EL), 2(4).

Osche, P.E., Castagnos, S., Napoli, A., and Naudet, Y. 2016. Walk the line: Toward an efficient user model for recommendations in museums. In Proceedings of 11th International Workshop on Semantic and Social Media Adaptation and Personalization.

Ryan-Collins, J., Stephens, L., Coote, A. 2013. The new wealth of time: How timebanking helps people build better public services. http://www.timebankingwales.org/userfiles/NEW

Smaldone, S., Lu, H., Shankar, P., Iftode, L. 2008. Roadspeak: Enabling voice chat on roadways using vehicular social networks. In Proceedings of SocialNets workshop, in conjunction with the IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications.

Srba, I., Bielikova, M., and Member, S. 2015. Dynamic group formation as an approach to collaborative learning support. IEEE Transactions on Learning Technologies, 8(2): 173-186.

Zhang, M. and Bockstedt, J. 2016. Complements and substitutes in product recommendations: The differential effects on consumers' willingness-to-pay. In Proceedings of Joint Workshop on Interfaces and Human Decision Making for Recommender Systems, in conjunction with the ACM International Conference on Recommender Systems.

Zhou, J., Sun, J., Athukorala, K., Wijekoon, D., and Ylianttila, M. 2012. Pervasive social computing: Augmenting five facets of human intelligence. Journal of Ambient Intelligence and Humanized Computing 3(2): 153–166.

## 5.2.   Content sharing

### 5.2.1. Survey

The increase in mobile data demand, coupled with the growing consumers' expectation to be connected in all places at all times, introduced significant research challenges in different areas of communication during the last decade. The widespread deployment of various network technologies, the increased adoption of multi-interface enabled devices (including support for ad-hoc networks) and the consolidation of peer-to-peer (P2P) content sharing schemes have enabled solutions in both solitary and collaborative forms (Lua et al., 2004). In the solitary form, the goal is to exploit any direct Internet connectivity on any of the available interfaces by distributing application data across them in order to achieve higher throughput, minimise energy consumption and/or minimise cost (e.g. of 3G/4G connections). In the collaborative form, the goal is to enable and incentivise mobile devices to utilise their neighbours' underutilised bandwidth in addition to their own direct Internet connections.

Utilising the mobile devices' interfaces in these two forms has been investigated over the past few years (Habak, Harras and Youssef, 2013). The focus has largely been on the solitary form developing a variety of multi-interface bandwidth aggregation solutions at different layers of the protocol stack (Higgins et al., 2010). Other attempts to utilize these interfaces in the collaborative form were also introduced (Sharma et al., 2004; Ananthanarayanan et al., 2007). These approaches either deal with a small-scale collaborative community managed by a single authority, or utilize proxy servers to handle and guarantee such collaboration. Overall, despite the fact that current smartphones, tablets, and other mobile devices are equipped with multiple

network interfaces, current operating systems mainly allow users to utilise only one interface at a time, or enable limited bandwidth sharing options through tethering. In other words, there has been a high deployment barrier and solutions have focused on bandwidth maximisation while not paying sufficient attention to energy and/or cost efficiency.

Wi-Fi Direct, initially called Wi-Fi P2P, is a Wi-Fi Alliance standard enabling devices to easily connect with each other without requiring a wireless access point. One advantage of this standard is the ability to connect devices even if they are from different manufacturers. Only one of the Wi-Fi devices needs to be compliant with Wi-Fi Direct to establish a P2P connection that transfers data directly between them with greatly reduced setup. The Android operating system has included support for this standard –through the so-called Wi-Fi P2P API (Google, 2016)– starting from version 4.0. This API sets a framework where apps can discover, connect, and communicate with other nearby devices (compliant or non-compliant with Wi-Fi Direct) across distances much larger than using Bluetooth, and even to create and manage groups for internal data sharing. Apple's mobile operating system, iOS, does not yet support Wi-Fi Direct and resorts to a proprietary solution called Bonjour (formerly, Rendezvous) instead (Apple, 2014). Regarding third-party initiatives to support direct communication among devices, it is interesting to mention the AllJoyn framework (AllSeen Alliance, 2013). This is an open source software project that makes it possible for apps to discover and communicate with one another, regardless of the underlying platform and with no need of Internet connection, only through Wi-Fi technology. The solution is currently available for Android and iOS devices, in addition to other computer platforms like Windows, OS X and Linux. In the AllJoyn framework, apps can create or join sessions, in peer-to-peer or multi-peer ways, i.e. it is possible to create groups for multi-user communications.

## 5.2.2. CrossCult applications

The main purpose of collaborative downloading and content sharing in CrossCult is about permitting device-to-device communication to distribute multimedia contents of common interest to the members of a sporadic social network who happen to be close enough to one another in the premises of a venue or around a city. As a secondary purpose, P2P could be also act to activate pre-loaded content in the device (app) to speed up communication. Applying the P2P philosophy, different devices would get different chunks of the content through Wi-Fi or 3G/4G connections, and then exchange those chunks through direct, inexpensive connections.

The collaborative downloading and content sharing features are intended to reduce the burden on the venues' Wi-Fi networks when multiple devices running the same app have to fetch the same contents. Instead of serving the whole files to each and every device, the infrastructure can ideally deliver them only once, and then the apps will engage in direct communications (over Wi-Fi, Bluetooth or other links) to get the full contents. In outdoors scenarios like those of pilot 4, this approach may help reduce the consumption of 3G/4G allowances.

## 5.2.3. Software, tools and algorithms

**Table 11: Content sharing software that will be possibly used in CrossCult.**

| Name | License | Short Description | Possible uses in CrossCult |
|---|---|---|---|
| **Wi-Fi P2P Direct** | Android | An Android API to establish direct Wi-Fi connections between nearby peers and to share files through that channel. | To exchange multimedia information among multiple devices, fetching data from other users reducing the required load to the venues' Wi-Fi infrastructures. |
| **AllJoyn** | CC Attribution 4.0 | A multi-platform open software framework for direct Wi-Fi communication among mobile and computer apps. | To exchange multimedia information among multiple devices, fetching data from other users reducing the required load to the venues' Wi-Fi infrastructures. |
| **ttorrent** | Apache Software License version 2.0 | A pure-Java implementation of the BitTorrent protocol, providing a tracker, a client and the related Torrent metainfo files creation and parsing capabilities. It can be embedded into apps. | To provide the logic to distribute the burden of downloading multimedia contents among different devices, and to exchange the chunks until getting to the complete contents. |

## 5.2.4. References

AllSeen Alliance. 2013. The AllJoyn framework. https://allseenalliance.org/framework

Ananthanarayanan, G., Padmanabhan, V.N., Thekkath, C.A. and Ravindranath, L. 2007. Collaborative downloading for multi-homed wireless devices. In Proceedings of 8th Workshop on Mobile Computing Systems and Applications (HotMobile), pp. 79–84.

Apple, Inc. 2014. Bonjour. https://support.apple.com/bonjour

Google, Inc. 2016. Wi-Fi Peer-to-peer API documentation. https://developer.android.com/guide/topics/connectivity/wifip2p.html

Habak, K., Harras, K.A. and Youssef, M. 2013. An optimal deployable bandwidth aggregation system. Computer Networks, 57(15): 3067–3080.

Higgins, B.D., Reda, A., Alperovich, T., Flinn, J., Giuli, T.J., Noble, B. and Watson, W. 2010. Intentional networking: Opportunistic exploitation of mobile network diversity. In Proceedings of ACM Mobicom, pp. 73–84.

Lua, E.K., Crowcroft, J., Pias, M., Sharma, R. and Lim, S. 2004. A survey and comparison of peer-to-peer overlay network schemes. IEEE Communications Survey and Tutorial, 7(2): 72-93.

Sharma, P., Lee, S.J., Brassil, J. and Shin, K.G. 2004. Handheld routers: Intelligent bandwidth aggregation for mobile collaborative communities. Technical report HPL-2003-37R1. http://www.hpl.hp.com/techreports/2003/HPL-2003-37R1.html

## 5.3.    Crowd management

### 5.3.1. Survey

The concept of crowd management refers to the way in which it is possible to manage the visitors of a physical space like a museum. In order to automate this process, it is necessary to know the user profiles and the context in which they are (Adomavicius, 2011). Once these elements are known, new opportunities emerge for crowd management (Naudet et al, 2014). One can for example quote the recommendation systems that exploit the user profiles to recommend items (in a museum, an item can be either an exhibit or a series of exhibits). Some dimensions of the user profile must be taken into effect in order to propose relevant items. Some of these have already been exploited in museum recommendation systems. One can for example quote the visiting style which defines the way a person visits a cultural site (Veron and Levasseur, 1989), or the implicit preferences (Bohnert et al., 2008), used to determine whether a person has more or less appreciated an item based on his/her past behaviour. However, a visit to the museum is a multi-dimensional event and other factors come into play to explain and understand visitor behaviour. Some ethnographic studies have identified and characterised some factors, such as the museum fatigue phenomenon that occurs when the visitor lowers the number of stops during the visit and is close to ending the visit. This phenomenon has been attributed to a long visit as well as a large number of exhibits seen (Davey, 2005). Detection and recognition of the concept of fatigue is not, to our knowledge, exploited in recommendation systems yet, but constitutes an important factor to explain the visitor behaviour.

Congestion in museums is also a known problem that has been studied for long. Studies have analysed visitor behaviour based on the time available for the visit and the exhibits they want to see (Yoshimura et al., 2014). To address this issue, computing the recommendations by both aiming to match the visitors' preferences and to reduce museum congestion points would temporarily divert visitors from the crowded areas. To this end, the work of Lykourentzou et al. (2013) proposes a multi-dimensional museum-visitor model, to simulate visitor behaviour inside the physical space of a museum. The model takes into account, on the one hand, factors related to the museum as a whole (particular layout, congestion capacity per room, etc.) and, on the other hand, visitor-specific elements (interest per exhibit, crowd tolerance, maximum available time, walking speed, time spent per exhibit, visiting style). The model is used to feed a crowd simulator (see table 12) that assesses the impact of different recommendation algorithms on the visitors' Quality of Experience.

### 5.3.2. CrossCult applications

CrossCult can take advantage of the technologies described in the survey for the purposes of controlling the trajectories of visitors in museums, which can be tested in the contexts of several pilots. Of particular importance is the National Gallery setting in the case of pilot 1, as it sees an impressive amount of visitors daily and thus is the most likely to benefit from congestion control. The recommendations provided by the CrossCult applications will allow visitors to avoid congested areas (and thus not adding to the congestion) in Pilot 1. Moreover, as the NG is the largest venue in the 3 pilots taking place indoors, the trajectories of visitors could inform a

user model which in turn can recommend rooms or paintings to new users based on predicted preferences based on the new visitors' early trajectories and the trajectories of past users.

### 5.3.3. Software, tools and algorithms

**Table 12: Crowd management software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|------|---------|-------------|--------------|
| **Crowd Simulator** | Proprietary (LIST) | Java-based software that simulates the movement of visitors inside the physical space of a museum based on a number of museum- and visitor-specific parameters (museum layout, visitor behaviour, visiting styles etc.) | To evaluate through simulation the effects of museum item recommendation algorithms on the Quality of Experience of visitors, and the museum-wide impacts that these algorithms have. |

### 5.3.4. References

Adomavicius, G., and Tuzhilin, A. 2011. Context-aware recommender systems. In Recommender systems handbook, pp. 217-253. Springer US.

Bohnert, F., Zukerman, I., Berkovsky, S., Baldwin, T., and Sonenberg, L. 2008. Using interest and transition models to predict visitor locations in museums. AI Communications, 21(2-3): 195-202.

Davey, G. (2005). What is museum fatigue. Visitor Studies Today, 8(3): 17-21.

Lykourentzou, I., Claude, X., Naudet, Y., Tobias, E., Antoniou, A. Lepouras, G., Vassilakis, C. (2013). Improving museum visitors' Quality of Experience through intelligent recommendations: A visiting style-based approach, Workshop Proceedings of the 9th International Conference on Intelligent Environments, J. A. Botia and D. Charitos (Eds.), IOS Press, pp.507 - 518.

Naudet, Y., and Lykourentzou, I. 2014. Personalisation in crowd systems. In Proceedings of 9th International Workshop on Semantic and Social Media Adaptation and Personalization, pp. 32-37.

Véron, E., and Levasseur, M. 1989. Ethnographie de l'exposition: l'espace, le corps et le sens. Centre Georges Pompidou, Bibliothèque publique d'information.

Yoshimura, Y., Sobolevsky, S., Ratti, C., Girardin, F., Carrascal, J. P., Blat, J., and Sinatra, R. 2014. An analysis of visitors' behavior in the Louvre Museum: A study using Bluetooth data. Environment and Planning B: Planning and Design, 41(6): 1113-1131.

# 6. Visualisation of associations and micro-augmentations

## 6.1.  Visualisation of digital cultural collections and associations

### 6.1.1. Survey

According to Card, Mackinlay and Shneiderman (1999), *information visualisation* can be defined as "*the use of computer supported interactive visual representation of data to amplify cognition*". Visualised cultural resource associations can reveal the various temporal, spatial, contextual, conceptual links between different cultural artefacts, their creators and associated events and lead to new ways of thinking about aspects of cognition as emergent properties of the interaction of people with cultural artefacts. Cultural heritage institutions are faced with content related challenges in order to publish their data collections: their content can be multi-format (text, images, audio, video, collection items, learning objects, etc.), multi-topical (art, history, artefacts, traditions, etc.), multi-lingual, multi-cultural and multi-targeted (targeted to both laymen and experts, various ages, etc.) (Heath, 2011). In the case of online museum collections, data associations build upon standard ontologies and vocabularies can have a beneficial impact on (i) exploring and navigating relationships, as richer semantics highlight the conceptual relationships between artefacts, and (ii) presentation-interaction, as they offer richer presentation possibilities in terms of browsing and navigation.

Case studies of museums that have implemented personalisation facilities to their websites show that understanding is stimulated when the systems use concepts familiar to the user (Bowen, 2004). Visuals help understanding by acting as a frame of reference or as a temporary storage area for human cognitive processes. By providing a larger working set for thinking and analysis they become external cognition aids (Kerren, 2008). Card, Mackinlay and Shneiderman (1999) list some key ways in which visuals can amplify cognition: increasing memory and processing resources available, reducing search for information, enhancing the recognition of patterns, enabling perceptual inference operations, using perceptual attention mechanisms for monitoring, and encoding information in a manageable medium.

When it comes to representing associations among multiple items, many of the aforementioned studies resort to graph representations, which appear recurrently in computer science to represent pairwise relationships among objects or concepts. Many important problems get clarified or even solved with the help of graph representations that improve their visualisation and open the door to graph theory techniques to analyse the relationships, detect issues and extract new knowledge (Larkin and Simon, 1987). Applications spread all along many other disciplines (e.g. Tufte 1998), ranging from technical scenarios (class hierarchies in software engineering, entity-relationship models in database systems, is-a relationships in knowledge representation, hyperlinks, etc.) to more general uses such as street layouts in maps or animal prey-predator relationships.

*Graph theory* (Tutte, 2001) is a mature discipline that helps answer very interesting questions related to a graph topology; for instance, whether there is a directed path between two given nodes, or between any pair of nodes, or to find the best route to go from one node to another. For CrossCult, the most appealing aspect of graph theory will be *graph drawing* (Gansner et al.

2002), a well-studied field that deals with methods for the generation of graph layouts, which can be very different for the same nodes and relations, affecting important properties like understandability, usability and aesthetics.

To this extent, we can distinguish –see (Cruz and Tamassia, 2015)– *algorithmic* approaches in which the layout of the graph is generated according to a pre-specified set of aesthetic criteria embodied in an algorithm as optimisation goals, or *declarative* ones, in which the layout is specified by a user-defined set of constraints. The former has the advantage of computational efficiency, but the latter is better at expressive power. A significant introduction to the field can be found in (di Battista et al., 1999). It is worth indicating that there are many different types of graphs and paradigms to tackle their visualisation (Didimo et al., 2014): hierarchical drawing, orthogonal drawing, overloaded orthogonal drawing, matrix-based representations, etc. Many of these algorithms are supported in several visualisation software packages, which readily provide support for the representation and management of connected graphs. Usually, as in the cases of Sigma.js, D3.js and Vis.js, these packages take the form of JavaScript libraries to be integrated and used in HTML pages, providing constructs to create connected graphs step by step or reading the data from a data structure or permanent storage. These libraries are heterogeneous in the support they provide, covering many different aspects in addition to simple graph representation, from input and interactivity management to graph animation. Some of the possible choices have recent maintenance activity and have proved to be compliant with the last generation of browsers.

## 6.1.2. CrossCult applications

One principal goal of the CrossCult project is to foster reflection about the common characteristics of distinct cultures in the European continent across different areas and ages. In this sense, visualisation allows people to identify connections about the assets (things, persons, traditions…) that represent such cultures in the venues; visualisation therefore can trigger reflection based on the information modalities it provides. This is a cross-pilot need, even appealing to some of the experts' tools offered by CrossCult. Visualisations of cultural resource associations can be complementary to textual descriptions for accessing and exploring cultural data and facilitate user's reflection and interpretation. Visualisations can help visitors, in a post-visit moment to discover the interconnectedness of digital cultural collections, enable presented information to be attuned to their interests and background, and therefore increase users' interest and engagement with both digital and physical collections.

Therefore, the envisaged functionalities involve the presentation of connected graphs associating objects, concepts, places, persons, topics etc. contained in the project knowledge bases. So, technologies will be needed to draw nodes and connections, attach labels, text, images, links etc. The tools selected and developed must permit the user interaction to display additional information (in the same screen or moving to a new one), to modify the graphs by rearranging the nodes or even adding new connections or attached data. Several open-source packages (see next section) seem to fulfil the project needs, especially knowing that the graphs to be displayed will not be large ones (presumably, 20-30 nodes on the visitors' apps, and up to 60-80 on the experts' tools). The computational cost of the algorithms is therefore not a decisive factor, just like their layout capabilities (all of them seem to be good enough displaying small graphs). The

ability to run smoothly in modest mobile devices and to enable interaction with nodes and connections will be important features. However, the cognitive effort of users to explore the graph in a busy setting as in most venues are the most decisive factors. Clarity of the graph presentations will be a priority, both to avoid forcing the users to reason about a large set of concepts and connections (which would make it hard to deliver or work out a single message, easy to be retained), and to avoid cluttering up the screens of their devices. In this sense, a natural choice to be studied will be *force-directed graphs* (Fruchteman and Reingold, 2001), which is a class of algorithms for drawing graphs in an aesthetically pleasant and intuitive way (allowing visitors to understand the data concisely, when there are other elements competing for their attention at a venue). Force-directed algorithms build graph visualisations as physical systems, assigning attractive and repulsive forces between nodes and edges to simulate their movement, seeking to minimise the total energy in the system.

Apart from graphs of nodes and connections, the CrossCult pilots need features to represent information of interactive timelines and maps. Those requirements can be fulfilled with HTML5, CSS3 and JavaScript technologies as well. Some of the libraries mentioned in Table 1 for graph drawing have extensions for drawing many different types of timelines. Embedding maps in web canvases is readily supported by Google Maps, Bing Maps and OpenStreetMaps, just to name a few.

## 6.1.3. Software, tools and algorithms

**Table 13: Visualisation software that will be possibly used in CrossCult.**

| Name | License | Short Description | Possible uses in CrossCult |
|---|---|---|---|
| **Sigma.js** | MIT License | Sigma.js is a JS library for drawing graphs, providing support for interactivity involving nodes and edges. | Drawing graphs of connections among concepts from the same or different venues in order to foster reflection about crosscutting aspects of history and culture in Europe. This is needed in the app of pilot 2, and in the experts' tools as regards curation of knowledge and design of pilot experiences. |
| **D3.js** | BSD license | A very powerful multi-purpose JS visualisation library, using HTML, SVG and CSS to render graphs based on a data-driven approach to DOM manipulation. | Drawing graphs of connections among concepts from the same or different venues in order to foster reflection about crosscutting aspects of history and culture in Europe. This is needed in the app of pilot 2, and in the experts' tools as regards curation of knowledge and design of pilot experiences. |
| **Vis.js** | MIT License | A browser-based visualisation library that supports many types of network/edge graphs, including auto-layout, auto-clustering, springy physics engine and animations. It also supports 2D/3D charts and timelines. | Drawing graphs of connections among concepts from the same or different venues in order to foster reflection about crosscutting aspects of history and culture in Europe. This is needed in the app of pilot 2, and in the experts' tools as regards curation of knowledge and design of pilot experiences. Presenting interactive timelines to show relevant temporal connections and context. Needed in pilot 2 and good for the experts' tools. |
| **SIMILE** | Open source | Free, open-source web widgets for | Presenting interactive timelines and |

| widgets | software under BSD license | data visualisations, maintained and improved over time by a community of open-source developers. Support to create web pages with advanced text search and filtering functionalities, with interactive maps, timelines, and other visualisations. | maps to show relevant spatio-temporal connections and context. Needed at least in pilots 2 and 4, and good for the experts' tools. |
|---|---|---|---|
| **Google Maps** | Commercial | A web mapping service developed by Google. It offers satellite imagery, street maps and other views, that can be embedded in web pages and enhanced with JavaScript programming to create mashups with other contents. | Presenting interactive maps to show relevant spatial connections and context. Needed at least in pilots 2 and 4, and good for the experts' tools. |
| **OpenStreetMap** | ODbL | A collaborative project to create a free editable map of the world. Interactive maps can be easily embedded in web pages and displayed in browsers with JavaScript support. | Presenting interactive maps to show relevant spatial connections and context. Needed at least in pilots 2 and 4, and good for the experts' tools. |

## 6.1.4. References

Bowen, J., Fantoni Filippini, S. 2004. Personalization and the web from a museum perspective. In Proceedings of Museums and the Web.

Card, S.K., Mackinlay, J.D. and Shneiderman, B. 1999. Readings in information visualization: using vision to think. Morgan Kaufmann Publisher, series in interactive technologies.

Larkin, J.H. and Simon, H.A. 1987. Why a diagram is (sometimes) worth ten thousand words. Cognitive science, 11(1): 65-100,

Tufte, E.R. 1998. Visual Explanations: Images and Quantities, Evidence and Narrative. Graphics Press.

Cruz, I. and Tamasssia, R. 2015. Graph drawing tutorial. Available online at http://www.fmf.uni-lj.si/~zitnik/dmii/graph_drawing_tutorial.pdf

di Battista, G., Eades, P., Tamassia, R. and Tollis, I.G. 1999. Graph drawing: Algorithms for the visualization of graphs. Pearson.

Didimo, W., Montecchiani, F., Pallas, E. and Tollis, I.G. 2014. How to visualize directed graphs: A user study. In Proceedings of 5th International Conference on Information, Intelligence, Systems and Applications.

Fruchteman, T. and Reingold, E. 2001. Graph drawing by force-directed placement. Software Practice and Experience, 21(11): 1129-1164.

Gansner, E.R., Koutsofios, E., North, S.C. and Vo, K.-P. 2002. A technique for drawing directed graphs. IEEE Transactions on Software Engineering, 19(3): 214-230.

Heath, T., Bizer, C. 2011. Linked Data: Evolving the web into a global data space. Morgan & Claypool, series Synthesis Lectures on the Semantic Web: Theory and Technology.

Kerren, A. 2008. Information visualization human-centered issues and perspectives.  Springer.

Tutte, W.T. 2001. Graph theory. Cambridge University Press.

## 6.2.  Augmented Reality and Micro-Augmentations

### 6.2.1. Survey

*Augmented Reality* (AR) is a live direct or indirect view of a physical, real-world environment whose elements are augmented (or supplemented) by computer-generated sensory input such as sound, video, graphics or GPS data (Augmented reality, 2016). Augmented reality has matured over the years and software tools (both libraries and end-user applications) that support it are now abundant. The user looks at the real world through a camera on a computing device; nowadays this is usually a handheld device such as a smartphone or tablet. The creator of the AR application has designated spots on the physical world where the generated content (augmentation) will be presented; when these spots come in the field of view of the user's camera, the AR application shows the relevant content.

The spots on the physical world where augmentations will be provided are usually marked with special markers, QR codes, etc. This helps the user locate them in space and turn his device towards them, and at the same time allows the application to position the augmentation more precisely. Markerless systems do exist, but in this case the user has to wander around hunting for the spots and the application must find their position by analyzing, in real time, the image captured at the device. If the content to be displayed is of visual form, it may pop-up when the user "sees" the designated point, aligned and scaled appropriately to the real object via computer vision (Craig, 2013); if it is auditory it may be played back automatically. The real-time constraint can be an important treason for choosing AR with markers or without, as the latter require more computational effort and time.

The micro-augmentations technology is a new tool, based on augmented reality, for presenting associations to end users in an unconventional format. Specifically, a micro-augmentation is a minimum meaningful stimulus provided to the user, of very short duration, carrying an emotional element and presenting the absolutely necessary information to capture attention. It creates an informative gap that leads to emotional arousal (Singer et al, 2009) and can increase learning motivation and visitor communications, thus solving real cultural heritage problems (e.g. isolating powers of technology, limited communication between visitors, low engagement of visitors, etc.).

The informative element provided by micro-augmentations is at a minimum level, allowing space for personal refection and meaning making. Micro-augmentations are location specific and they support direct visitor-exhibit inner dialogue. The stimuli are also at the minimum conscious perception level, moving between perception and intuition (a low sound may cause a user to wonder whether s/he heard it or not). This uncertainty leads to emotional arousal (Singer et al, 2009).

In addition, micro-augmentations are based on the careful selection of stimuli. Sounds and visual stimuli are selected not solely on the basis of creating an ambience but attempting to provide the visitor with meaningful stimulation that could trigger different cognitive and emotional processes or just recover visitors' attention or driving it to the desired POI (as discussed in Section 5.3). For example, specific sounds are carefully selected (e.g. tonalities to be used are As-dur, A-dur, H-dur, E-mol), based on their relations to particular emotional states (happiness, great energy, dreaminess), using findings from psychological research (Sapozhnikova and Taymanov, 2009). Thus, micro-augmentations also investigate the role of emotions in the overall museum experience, applying knowledge from cognitive psychology (e.g. Sapozhnikova and Taymanov, 2009). Emotions can trigger a number of cognitive processes, like motivation for action (Pekrun et al, 2011), positive learning outcomes (Izard, 1991), increase in satisfaction levels (Um et al., 2007), since they involve affective, cognitive, physiological, motivational, and expressive components (Scherer, 2009). The recent study of (Antoniou et al., 2015) showed increased user engagement and increased positive emotions in museum settings.

Micro-augmentation is not a new technological component, but it rather uses existing technology in a novel way and it is a new method of content delivery. In this light, micro-augmentations could use visual and/or audio stimuli. Micro-augmentation tool are implemented as a software module (library) that retrieves the required triggers about micro-augmentation playback from the back-end, and invokes the relevant modules of the mobile app for audio/visual info presentation to the user.

## 6.2.2. CrossCult applications

In the CrossCult project, 3D visuals will be used in two different ways in the mobile apps: (a) to present content in augmented and virtual reality contexts and (b) to display 3D models. In all pilots of the project, the mobile app will display cultural items related to items at the site the users are currently visiting. Augmented Reality, as presented above, can effectively help to better put the presented material in context and to enhance user engagement; therefore, augmented reality techniques will be used, when and where it will be deemed appropriate, to present interconnected material. Virtual reality techniques (rendering complete worlds in immersive experiences) may also be applied, given the recent technological advancements and availability of VR headsets, as long as they do not disrupt the flow of the visit, and only if pilots require it.

3D models of cultural items, such as pots, statues, etc., may (subject to availability of the models) be displayed to the users when they explore cultural associations through the CrossCult mobile app if pilots require their use. The advantage of using 3D models of the objects instead of simple, static images has been outlined above. Furthermore, the issues related to bandwidth needs and processing power for their rendering are hardly relevant for our tasks: on one hand, it is neither technically difficult nor expensive to provide good and high-speed WiFi coverage of indoor or outdoor sites (which will, in any case, be also required for our app to function properly); on the other hand, current smartphones possess powerful Graphics Processing Units (GPUs) and implement 3D graphics APIs (OpenGL), so that rendering modest size 3D models will not be a problem.

## 6.2.3. Software, tools and algorithms

**Table 14: Augmented and Virtual Reality software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|---|---|---|---|
| **ARToolkit** | LGPL 3.0. | A full-featured AR Toolkit. It offers robust tracking and calibration algorithms, using both markers and natural features. It supports the Unity3D and OpenSceneGraph graphics engines. It is available on all major mobile and desktop platforms. | Display related assets from associated venues in an augmented-reality context. |
| **Wikitude Studio** | SaaS or License | Set of products and plugins to develop 3D AR experiences for mobile. It has a programmers' SDK and plugins for HTML5 app builders frameworks | 2D/3D augmented reality for Pilots and CCFE |
| **Alvar** | VTT Propietary - Licensing | Suite of products and SDKs for creating augmented reality applications. ALVAR provides leading tracking technology that can be leveraged to add value to products and services, solve industrial problems and enable next generation tools | 3D markerless agumented reality for Pilots and CC FE |
| **DroidAR** | GPL 3.0 / Commercial | DroidAR is a framework for Augmented Reality on Android. Location based AR and marker based AR are both possible. Its development seems to have stagnated. | Display related assets from associated venues in an augmented-reality context. |
| **OpenSceneGraph** | Open source, OpenSceneGraph Public License (based on the LGPL) | The OpenSceneGraph is open source, real-time graphics middle-ware used by application developers in fields that range from visual simulation (flight, marine, vehicle, space simulator) to virtual and augmented reality, to medical and scientific visualisation, to education and games. | Advanced 3D content display. |
| **Unity3D** | Commercial / Free for personal or educational use | An advanced game engine with support for all current desktop, mobile and game console platforms | Advanced 3D content display. |

## 6.2.4. References

Augmented                                reality.                                2016. https://en.wikipedia.org/w/index.php?title=Augmented_reality&oldid=741082168   Retrieved   09/20, 2016

Craig, A.B. 2013. Understanding augmented reality; Concepts and applications. Morgan-Kaufmann.

Antoniou, A., O'Brien, J., Barnes, A., Bardon, T., and Virk, D. 2015. Micro-augmentations: Situated calibration of a novel non-tactile, peripheral museum technology. In Proceedings of 19th Panhellenic Conference on Informatics, pp. 229-234.

Izard, C.E. 1991. The psychology of emotions. Plenum Press, New York.

Pekrun, R., Goetz, T., Frenzel, A.C., Barchfeld, P., and Perry, R.P., 2011. Measuring emotions in students' learning and performance: The Achievement Emotions Questionnaire (AEQ). Contemporary Educational Psychology, 36: 36-48.

Ritchie, G. 2007. Some empirical criteria for attributing creativity to a computer program. Minds and Machines, 17(1): 67-99.

Sapozhnikova, K., and Taymanov, R. 2009. Multiparametric measurements of emotions. In Proceedings of XIX IMEKO World Congress Fundamental and Applied Metrology, pp. 1180-1185.

Scherer, K. 2009. The dynamic architecture of emotion: Evidence for the component process model. Cognition and Emotion, 23: 1307–1351.

Singer, T., Critchley, H.D., and Preuschoff, K. 2009. A common role of insula in feelings, empathy and uncertainty. Trends in Cognitive Sciences, 13(8): 334-340.

Um, E.R., Song, H., and Plass, J. 2007. The effect of positive emotions on multimedia learning. In Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications, pp. 4176-4185.

# 7. Integration and alpha testing

## 7.1.  Collaborative software development

### 7.1.1. Survey

In 1968, Melvin E. Conway concluded a paper about designing system with the following observation: "*Any organisation that designs a system (defined broadly) will produce a design whose structure is a copy of the organisation's communication structure*" (Conway, 1968). This observation, known as *the Conway lay*, has been revisited many times in the realm of software development, which has collaboration as a cornerstone: virtually all software development requires collaboration among developers within and outside their project teams, to achieve a common objective. It was early shown that about 70% of a software engineer's time is spent on collaborative activities (Vessey and Sravanapudi, 1995), and the tendencies in the modern global market do not call for a significant revision of that figure (Treude, Storey and Weber, 2009). Indeed, collaboration in software development has been extensively studied by researchers in the fields of *Software Engineering* and *Computer Supported Cooperative Work* (CSCW) since the 1980s, producing a wide range of collaborative development tools (Mistrík et al., 2010).

Research has produced a host of tools, each typically focusing on a different aspect of collaboration, as well as a range of methodologies applicable in different settings. Several classification frameworks exist that can be used to classify those solutions: some provide a detailed taxonomy to compare tools in a particular area (Conradi and Westfechtel, 1998), some classify tools based on the functionality of the tools (Grudin, 1994), some classify tools based on the high-level approach to collaboration that the tools take (Al-Ani et al., 2008) and so on.

As the basic trends in modern software development, it is possible to highlight the following: (i) *software reuse*, considered as the process of creating software systems from existing software rather than building them from scratch, which often helps improve software quality and development productivity (Leach, 2013); (ii) *agile methodologies*, defining workflows that seek to reduce overheads in the software process (e.g. by limiting documentation) and to be able to respond quickly to changing requirements without excessive rework, on the basis of interleaving specification, design, implementation and testing tasks with continuous stakeholder involvement (Cockburn and Williams, 2003); (iii) *refactoring*, looking for possible software improvements and making them even where there is no immediate need to do so, seeking to improve the understandability of the software, reduce the need for documentation, and ease further development (Kim et al., 2012); (iv) *continuous integration*, a development practice that requires developers to integrate code often and to run automated tests after each check-in as a means to detect problems early (Olsson and Bosch, 2014); and (v) *agile project management*, adapted to incremental development as per the guidelines of methodologies like Scrum (Rising and Janoff, 2000).

## 7.1.2. CrossCult applications

As stated by Conway's law, the quality of the CrossCult platform will be directly correlated to the quality of the processes in place to integrate the different technologies delivered by WP3 tasks. The challenge will be to coordinate the efforts of geographically-distributed teams, working on specific technologies rather than specific functionalities. The integration should rely on industry-grade processes and tools on one side, but also on adapted architectural choices. The overall architecture of the platform should support "by design" the way all the partners will work, e.g. using different languages, different technologies, working at different rhythms on specific, well identified parts of the platform. The idea is to keep things small, simple and with clear use cases to maintain maximum flexibility and reach quick and regular outcomes.

A recent architectural style that seems to fit the CrossCult requirements is the *microservice* style, put forward in (Fowler and Lewis, 2014) as *"an approach for developing a single application as a suite of small services, each running in its own process and communicating with lightweight mechanisms, often an HTTP resource API. These services are built around business capabilities and independently deployable by fully automated deployment machinery. There is a bare minimum of centralised management of these services, which may be written in different programming languages and use different data storage technologies."* In that sense, the CrossCult platform has been posed as a suite of microservices instead of a big monolithic application. Each technology developed in WP3 will be integrated inside a task specific microservice that will have been designed to support a well-identified requirement of the platform. The main point of such an architecture is to lower the overall complexity by breaking down a system into smaller, more easily manageable components that can be handled individually by a development team. This fosters iterative developments that can meet regular, reachable objectives, which is in the line of agile development methodologies.

To support the development process, and as previously described in deliverable D1.1, it has been decided to use Gitlab as a collaborative development platform. A project-specific instance of Gitlab has been set up and is hosted by LIST. This platform provides tools to maintain version control of source code (based on GIT), and facilities to organise the different code repositories (through groups and projects) and to support discussions and developer interactions via the issue tracker. All the software produced in the framework of the project should be stored and managed through the CrossCult Gitlab. The documentation of all the modules developed can be hosted on the platform either through the built-in wiki or as versioned files along with code. This documentation will be delivered along with the software modules as part of D3.2 and D3.4.

To facilitate the integration of the different modules, the project will rely on a continuous integration process. In a few words, the idea is to avoid the pitfalls of a strict waterfall process by introducing continuous testing and verification along the implementation phase. This will allow detecting any misalignment or critical issue early. The details of this process will be discussed more at length as part of deliverable D3.3.

## 7.1.3. Software, tools and algorithms

**Table 15: Collaborative development software that will be possibly used in CrossCult.**

| Name | License | Description | Possible use |
|---|---|---|---|
| **GitLab** | Copyright 2011-2016 GitLab B.V. | A collaborative development platform that can be installed on premise and provide most of the tools and facilities to support software development. | The main code repository. The platform will also be used to store documentation and for reporting/track bugs. |
| **gitlab-ci-multi-runner** | MIT | A Continuous Integration tool that integrates with Gitlab. | Use it to set up Continuous Integration for all the code repositories of the project. |
| **Sonarqube** | LGPL v3 | A code quality assessment platform that allows analysing code source and generating quality reports based on best practices and user provided coding rules. | To be used for code quality assessment. |
| **Spring Boot** | Apache 2.0 | A Java framework that can be used as a microservice chassis (Richardson 2014), providing embedded features such as log management, authorisation management and easy configuration facilities. | To be used as a microservice chassis for building the platform services. |
| **Swagger** | Apache 2.0 | Swagger is a simple yet powerful representation of your RESTful API with support in almost every modern programming language and deployment environment. It includes interactive documentation, client SDK generation and discoverability. | Use to document and formalize services API in a common standard to facilitate interoperability between services and apps. |

## 7.1.4. References

Al-Ani, B., Trainer, E., Ripley, R., Sarma, A., Van Der Hoek, A., Redmiles, D. 2008. Continuous coordination within the context of cooperative and human aspects of software engineering. In Proceedings of 2008 International Workshop on Cooperative and Human Aspects of Software Engineering, pp. 1-4.

Cockburn, A., and Williams, L. 2003. Agile software development: It's about feedback and change. Computer 36(6): 39-43.

Conradi, R., and Westfechtel, B. 1998. Version models for software configuration management. ACM Computing Surveys, 30(2): 232-282.

Conway, M.E. 1968. How do committees invent? Datamation magazine. [Accessed 09 2016]

Fowler, M., and Lewis, J. 2014. Microservices, a definition of this new architectural term. http://martinfowler.com/articles/microservices.html [Accessed 09 2016]

Grudin, J. 1994. CSCW: History and focus. IEEE Computer, 27(5): 19-27.

Kim, M., Zimmermann, T., and Nagappan, N. 2012. A field study of refactoring challenges and benefits. In Proceedings of the ACM SIGSOFT 20th International Symposium on the Foundations of Software Engineering.

Leach, R.J. 2013. Software reuse: Methods, models, costs. Ronald J. Leach, 2nd edition.

Mistrík, I., Grundy, J., van der Hoek, A., abd Whitehead, J. (Eds.) 2010. Collaborative Software Engineering. Springer.

Olsson, H.H., and Bosch, J. 2014. Towards agile and beyond: An empirical account on the challenges involved when advancing software development practices. Lecture Notes in Business Information Processing, 179: 327-335.

Richardson, C. 2014. Pattern: Microservice chassis. http://microservices.io/patterns/microservice-chassis.html [Accessed 09 2016]

Rising, L., and Janoff, N.S. 2000. The Scrum software development process for small teams. IEEE software, 17(4): 26-32.

Treude, C., Storey, M.A., and Weber, J. 2009. Empirical studies on collaboration in software development: A systematic literature review. Technical Report DCS-331-IR, Department of Computer Science, University of Victoria.

Vessey, I., Sravanapudi, A.P. 1995. Case tools as collaborative support technologies. Communications of the ACM, 38: 83-95.