International Review of Research in Open and Distributed Learning

A MOOC on Approaches to Machine Translation

Marta Ruiz Costa-jussà, Lluís Formiga, Oriol Torrillas, Jordi Petit and José Adrián Rodríguez Fonollosa

Cite this article

Ruiz Costa-jussà, M., Formiga, L., Torrillas, O., Petit, J. & Rodríguez Fonollosa, J. (2015). A MOOC on Approaches to Machine Translation. *International Review of Research in Open and Distributed Learning*, *16*(6), 174–205. https://doi.org/10.19173/irrodl.v16i6.2145

Article abstract

This paper describes the design, development and analysis of a MOOC entitled "Approaches to Machine Translation: rule-based, statistical and hybrid" providing lessons learnt on conclusions to be take into account in the future. The course was developed within a Canvas platform, used by recognized European universities. The course contains video-lectures, quizzes and laboratory assignments. Evaluation is done across on-line quizzes, programming assignments (PAs) evaluated by means of a specific code evaluation and peer-to-peer strategies. This MOOC allows to introduce people from various areas to the Machine Translation theory and practice. It also allows to internationally publisize different tools developed at the Universitat Politècnica de Catalunya.

November – 2015

A MOOC on Approaches to Machine Translation

Marta R. Costa-jussà⬚, Lluís Formiga✦, Oriol Torrillas, Jordi Petit and José A. R. Fonollosa

⬚Instituto Politécnico Nacional, Mexico,

✦Verbio Technologies S.L., Barcelona

Universitat Politècnica de Catalunya, Barcelona

Abstract

This paper describes the design, development, and analysis of a MOOC entitled "Approaches to Machine Translation: Rule-based, statistical and hybrid", and provides lessons learned and conclusions to be taken into account in the future. The course was developed within the Canvas platform, used by recognized European universities. It contains video-lectures, quizzes, and laboratory assignments. Evaluation was through on-line quizzes, programming assignments assessed by means of a specific code evaluation, and peer-to-peer strategies. This MOOC allowed people from various fields to be introduced to the theory and practice of Machine Translation. It also enabled us to internationally publicize various tools developed at the Universitat Politècnica de Catalunya.

Introduction

Several years after the first Massive Open Online Course (MOOC) by David Wiley in 2007[1], this type of course has received a great deal of attention from academics as a way to disseminate different levels of knowledge with the contribution of the participants. Among their various advantages, MOOCs require low levels of resources in relation to their potential international impact.

---

[1] http://creativecommons.org/tag/massive-open-online-course

In the context of Natural Language Processing (NLP), Machine Translation (MT) brings together varied communities from linguists to engineers. From the Universitat Politècnica de Catalunya (UPC), with more than 12 years of experience in the topic, we launched a MOOC that provides an overview of the basis for the most popular MT systems. Our main motivations for developing this course are listed as follows:

- To improve the content and digital resources of MT.

- To adapt an MT course to an international public.

- To propose a course on a complete platform that is used by large European universities.

- To use and publicize the open-source NLP tools developed at the UPC that are useful for the MT community.

Although the MOOC we are describing is very specific to MT, it was well received, with over 772 students registering from all around the globe. The most relevant characteristic of the MOOC presented is its development under a low-resource scenario and its far-reaching marketing strategy using in-house low-budget tools. The MOOC was set up on the well-known Canvas platform, but customized to UPC requirements. In this sense, this paper extends our ongoing previous work (Costa-jussà, Formiga, Petit & Fonollosa, 2014) and describes the complete course, focusing on how the low-resource MOOC on the topic of MT has been developed, together with a deep analysis of the output of the course mainly in terms of participation. Therefore, this paper can be useful to professors and/or researchers that want to develop a MOOC on a fresh topic and independently from the big platforms. Among the different ideas that the reader can use, we provide details of how to: constitute a multidisciplinary group of people to develop the course; promote and motivate the course; proceed step-by-step for recording and preparing evaluation materials; and design and analyze the pre- and post-course surveys.

The rest of this paper is organized as follows. Section 2 depicts a brief description in both courses of MT with open material and already existing MOOCs under similar fields. Section 3 presents the course topic and contents. Section 4 describes the procedure for recording the video lectures and the promotional video, the functionalities of the Canvas platform, the MOOC team and how the course was promoted. Section 5 includes an analysis of the students including a demographic description, as well as a study into their background and experience of MOOCs. Section 6 gives details about participation in the course, such as how many people watched video-lectures, performed activities, and participated in the discussions; the section also describes the feedback received from students, and finally, it defines the correlation between participation and feedback, among other interesting issues. The last section concludes the paper with some closing remarks.

Related Work

MOOC has become a strong stakeholder in professional development; however, up until now, there was no evidence of any MOOC explicitly related MT. To the best of our knowledge, the most

similar MOOC is a course oriented to NLP in 2012 by Coursera[2], a pioneering platform for such courses.

Moreover, as will be detailed in Section 3, the MT community has put a lot of educational effort with the publication of several books, on-line slides, open source tools, tutorials, workshops, and self-assessment coding tools. Nevertheless, there are fundamental characteristics defining the particularities of this subject; MT itself lies between several autonomous fields such as engineering and linguistics, hence it emerges as a frontier for artificial intelligence. The profile of MT experts require both a strong background in engineering due to its statistical facet and major notions from the field of linguistics. However, undergraduate schools today teach the intricacies of the two combined disciplines, a field known as computational linguistics. Moreover, there is a consensus on the best way to teach and make comprehensible the dynamics of MT is a methodology of Problem-Based Learning (PBL). In that sense, moving from a teaching methodology with strong practical requirements to a massive environment becomes an added challenge to the MOOC presented. It is also important to note the peculiarities of our working scenario:

> 1. Deploy an initial MOOC based on expertise (not introductory) that combines students from linguistics and engineering without a strong bias to both disciplines.

> 2. Convert the teaching material that has traditionally been delivered in classroom, long-run format at a university - sessions of 2 hours without much on-line support - to a 10-minutes online-pill format lessons.

> 3. Infer a PBL methodology where students cooperate with their work under multidisciplinary teams and therefore encouraging engineers to help linguists and vice-versa.

Next, we review some of the related work concerning these points of the scenario. Despite not all of that work has been considered into MOOC design; some parts are enumerated for their importance to the field.

One of the first thorough studies about Open and Distance Learning (ODL) and linguistics was provided by Godwin-Jones in 2012, where the author introduced the trend in Western countries today to instruct language learning in higher education. Normally, the use of the Web was included as an additional resource in face-to-face classes, as an equal partner in instructional delivery in hybrid courses, or as the primary teaching and learning environment in distance learning. However, our course does not focus on teaching a specific language to humans, but on how machines can learn to translate. Nevertheless, some of his practical ideas helped us to design the working assignments for our course. In the same direction, (Abeer & Miri, 2014) examined students' preferences and views about learning in MOOCs and specified the criteria that make MOOCs a constructive learning environment based in its learning-by-doing philosophy. Another

---

2 https://class.coursera.org/nlp/lecture

topic beyond the intention of our course but on the basis of computational linguistics, is the increasing importance of emotional automatic analysis concerning students' messages in forums, discussions, and opinions of the course. To that effect, Yang, Liu, Liu, Min, and Meng (2014) selected relevant features from original feature, train-based classifiers to build a dynamic and robust ensemble. The experimental results on the real-life moods posts demonstrated that their methodology outperformed conventional multi-view semi-supervised emotion recognition methods. With the same intention, but under the ontology perspective, Gil, Virgili-Gomá, García, and Mason (2015) applied ontologies in the context of Emotional Common Sense, collaboratively collecting emotion common sense and modelling it with an ontology named EmotionsOnto. Their experiments were conducted to automatically measure users' emotional states using Brain Computer Interfaces. Moreover, under the social perspective Shen and Kuo (2015) depicted how many MOOCs have adopted social media tools for large student audiences to co-create knowledge and engage in collective learning processes, by adopting various social media mining approaches to investigate Twitter messages related to MOOC learning. In short, the findings pertaining social media mining afforded a holistic understanding of MOOC trends, public sentiment toward MOOC learning, and the influencers of MOOC-related re-tweets.

Thanks to the rapid development in technology, which can be a common ground for many fields, interdisciplinary studies have gained importance resulting in an integration of different concepts, into ODL The field of ODL has been growing requiring research in theory and their application (Aydemir, Özkeskin, & Akkurt, 2015). A strong evidence of this growth is given in the research of Muñoz-Merino, Ruipérez-Valiente, Alario-Hoyos, Pérez-Sanagustín and Delgado Kloos, C. (2015) where Precise Effectiveness Strategy (PES) is presented. PES defines precise metrics allowing the calculation of the effectiveness of students when interacting with educational resources and activities within MOOCs and Small Private On-line Courses (SPOCs). It has to be taken into account the particular aspects of the learning context. PES visualizations can help teachers make quick and informed decisions in their case study, enabling the whole comparison of a large number of students at a glance, and a quick comparison of the four SPOCs divided by videos and exercises. Another important aspect to consider when designing a MOOC is the on-line social behaviour, especially when transferring existing syllabus dynamics from classrooms to on-line platforms. In that sense, Kolp and Wautelet, (2014) raised the strong link between ODL and social networks behaviour with the clear intention to disaggregate the analysis of MOOCs. More precisely, they overviewed a set of human organizational styles from literature's theory. These theories were applied in an original manner to the representation of collaborative learning software systems. By using socially driven software development methodologies like Tropos, a new software system could be modelled as a set of collaborating actors pursuing collective and individual goals. That team effort, in addition to the individual effort, was necessary to obtain a competitive and professional solution (Harms, 2015) under corporate environments. Therefore, Self-Regulated Learning (SRL) has shown to be an effective way to learn in entrepreneurial projects at the level of the individual student, demonstrating the effects of SRL, team learning, and psychological safety on various assessment types in the context of an entrepreneurship class. That evidence becomes a key factor in our MOOC as most of the human translation professionals

work as self-employed professional liberals allowing inferring the potential of Computer Aided Translation (CAT) into their own communities.

Problem Based Learning (PBL) is an approach to learning where the syllabus is designed with problem scenarios central to student learning in each curricular component (modules/units). Terms such as "computer-mediated PBL" and "online PBL" have been used to define forms of PBL that utilise computers in some way. Despite this, many of the concerns raised by delegates at conferences around the world (Savin-Baden & Wilkie, 2006): include whether online PBL can:

- affect the existence of face-to-face PBL since online PBL will be seen as being more cost effective.

- destroy some of the original aims of PBL since some forms of online PBL tend to focus on solving narrowly defined problems that fail to encourage students to be independent inquirers who own their learning.

- reduce the impact of learning in teams, in terms of students learning to work through team difficulties and conflicts in the way required by face-to-face PBL.

Some forms of online learning are difficult to marry with types of PBL that seek to provide opportunities for the students to challenge, evaluate, and interrogate models of action, knowledge, reasoning, and reflection, such as PBL for critical contestability (Savin-Baden, 2000). Rendas and Gamboa (1999) introduced a computer simulation that was designed for PBL in order to motivate learning and structure knowledge. In the following sections, we detail how we tailored the content and activities of the course to a PBL methodology.

Course Topic and Content

The aim of the course was to provide a broad view of MT. Given that MT is a multidisciplinary field involving linguists, computer scientists, engineers, and computer technicians, this course was intended for all these profiles. The course reviewed the most important paradigms of MT including statistical, rule-based, and hybrid translation, giving an overview of state-of-the-art techniques and algorithms, and it outlined the most challenging problems in the field. In addition, an overview of several evaluation approaches was presented: automatic measures; human perception evaluation; and human linguistic evaluation, with emphasis on the advantages and disadvantages of each method. The course underlined how statistics and linguistics are combined in both development and evaluation and was organized over five weeks with the addition of two extra weeks. Table 1 shows the distribution of topics per week; the last two weeks were dedicated to the development of activities and the final exam.

Table 1

Course Organization

Week	Topic
1	Introduction and rule-based MT
2	Alignment and phrase extraction
3	Reordering, language modeling, tuning and decoding
4	Evaluation and toolkits
5	Hybrid MT and MT use case study

From the theoretical point of view, at the end of the course the student was expected to understand the theory behind rule-based, statistical, and hybrid MT systems. The students were asked to pass several theoretical tests to demonstrate their skills. From a practical point of view, at the end of the course the student should be able to build a rule- and a statistical-based MT system. Students were also asked to build parts of the systems during the course.

Key references for the course included the Statistical Machine Translation book (Koehn, 2010) and the survey of Linguistics applied to Statistical Machine Translation (Costa-Jussà & Farrús, 2014). Other references included slides from Adam Lopez[3] and corresponding activities (Lopez et al., 2013). Finally, it should be mentioned that the course was compiled solely from the teachers' experience and in cooperation with other top researchers while running courses on the same topic in the RUSSIR 2012 Summer School[4], Universidade de São Paulo, Universitat Politècnica de Catalunya, ESSLLI 2013 Summer School[5] as well as giving talks and tutorials around the world including ESWC 2013[6] and NAACL 2013 (Costa-jussà, 2013).

The course activities constituted the main differential linchpin of the MOOC while they contributed to its work-flow identity. In addition, they made up the set of tools needed in order to obtain qualitative and quantitative metrics to further understand what was taking place within the student learning process.

However, when designing activities, it was important to take into account the fact that not all MOOC content may be inclusive to all participants. Specifically, people coming from areas such as linguistics, humanities, fine-arts, or translation and interpretation may not have a solid mathematical background. Thus, two non-exclusionary profiles were established to achieve the certificate of accomplishment based on the syllabus. The *mt-developer* profile was oriented towards candidates with a strong mathematical background, and the *mt-manager* profile oriented towards candidates with a linguistic background and an interest in abstractly understanding several concepts and being able to interact with them.

---

[3] http://mt-class.org
[4] http://videolectures.net/russir2012_yaroslavl/
[5] http://esslli2013.de/schedule/
[6] http://webmining.barcelonamedia.org/lambert/crosslingmind/eswc2013-tutorial/

From the first day, teachers and collaborators aimed to identify the roles of leader and follower among the participants based on different variables: i) participation in the forums; ii) test grading; iii) behavioral patterns both in taking the lessons in the form of slides and video lectures. Leaders were publicly reputed, and their motivation was augmented using techniques learned from social networks through authority, popularity, and typical metrics such as like, dislike, mentions, and badges.

Nevertheless, we set the focus of collective work beyond the scope of a forum or informal social interaction, including the activities. The working plan was not limited to programming assignments and weekly quizzes, but also explored other skills the students may have in order to encourage peer cooperation and make the learning process become more of a shared journey than an individual achievement.

For this reason, we designed the scientific assignment to require student complicity; students must cooperate to solve several challenges that they may encounter during the course by mutually correcting scientific-paper reviews or by competitively brainstorming ideas to guide the MT research community. In addition, we worked from a PBL perspective. The purpose of the programming assignment was not only to provide a challenge for the *mt-advanced* candidates through programming the core lines of source code but also for the *mt-manager* students, to be able to build a bash script for training a completely ad-hoc MT-system based on freely available tools.

PBL and surfing the internet share similar qualities, for example the process of learning in PBL teams is interactive, non-sequential, random, and often seems rather chaotic. Furthermore, as with students undertaking face-to-face PBL, those students doing online PBL seem to need less facilitator support and guidance as they become more familiar and skilled with managing the learning approach. For example, in face-to-face PBL, students have often reported feeling out of control, fragmented, and ill at ease because the experience of learning through PBL is markedly different from prior learning experiences (see for example Savin-Baden, 2000). Online PBL offers the students greater autonomy to learn for themselves and opportunities to develop independence in inquiry, opportunities many of them may not have encountered in lecture-based learning. To accomplish PBL, the programming assignments were automatically evaluated using the Jutge.org platform (Petit, Giménez, & Roura, 2012). Jutge.org is an open access educational on-line programming judge that students can use to try to solve a broad range of programming problems using several programming languages. The verdict on their solutions is computed using exhaustive test sets run under time, memory, and security restrictions. Further details of this platform are explained in Section 5. The detailed structure of the activities, following a PBL design, was:

- Weekly Quizzes - Trial-and-error tests with random questions and answers to measure the how well the student had understood content of the syllabus.

- Elementary Programming Assignments - **T**hese were based on bash scripting evaluating the appropriate calls to freely available toolkits such as SRILM, Moses or Asiya (Giménez

& Màrquez, 2010) in terms of argument passing and coherence. The student not only learns how to build a script that performs all the parts of the process but they can also check the correctness of each outcome by uploading it to the server and make sure they understand it by means of a simple quiz.

- Advanced Programming Assignment - An automatic software evaluator based on a DreaMT of a student's code.

- Scientific Assignment - The students had to review a state-of-the-art MT paper and provide some dissemination alongside future research ideas. They were peer-evaluated.

- Final Exam: -The final exam gave the student a global perspective. However, it only had a marginal effect on the final grading.

It is worth mentioning that the majority of the course's success is due to the flexibility and re-usability of the Canvas platform thanks to its LTI[7] interfaces and applications. This is explained in more detail in Section 5.

Building the MOOC: Technical and Practical Procedures

This section describes several details on the procedures of creating the MOOC. The section includes details on recording, platform, team and promotion.

Recording Set

The video recording sessions for the MOOC posed various challenges. The most important aspect involved grabbing the student's attention in order to transfer the course content efficiently. A MOOC has certain characteristics that differ from a normal class-based lecture. For example, it is necessary to avoid long sessions involving a great deal of mathematical or conceptual development. For this reason the communication is more direct and requires a delivery structured in bite-sized chunks where each chunk transmits very few concepts but is very well prepared.

For this reason the aesthetics and means of obtaining audio and video were much more important than in a master class at the faculty. It is important to bear in mind that the overall quality of the video session was evaluated as much as its ability to transmit specific knowledge in an affordable way, as was the ability to unlock a student from a jamming situation, which could be caused by a misunderstanding of the slides or the fact that they are not self-explanatory enough.

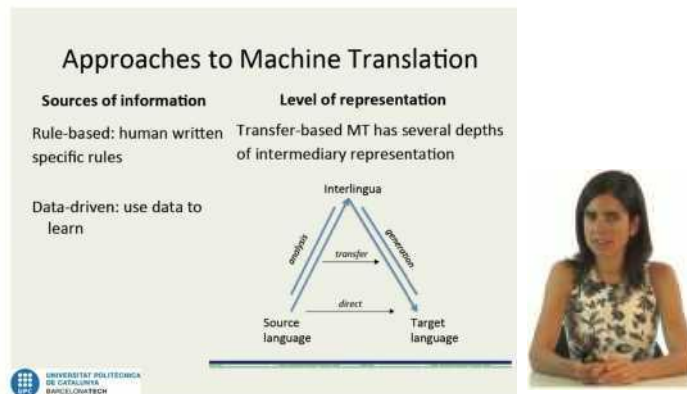[7] http://www.imsglobal.org/toolsinteroperability2.cfm

Figure 1. Screen showing a recording of the teacher and the presentation slides

In this section we detail the design decisions made for the MOOC video recordings in order to increase the success of the course. In order to make these decisions, we followed advises from well-known sources[8]. First the maximum length of the videos was set as between 10 and 15 minutes. Moreover, sound quality was prioritized using a directional lapel microphone: this was the main means of communication between teacher and student. The sessions were held on a minimalist set: a flat white background and a waist up close-up of the sitting teacher, capturing her facial expressions and body language. The main display for the student consisted of the presentation slides, occupying most of the area of the screen, with the teacher in a discreet corner, and a corporate background (see Figure 1).

On the recording set, the teacher could view the presentation in two places, first on a touch screen device on the table such as a Wacom tablet or iPad[9] that allowed them to delete, write, annotate, or make clarifications of any kind on the student's screen in order to focus on specific aspects. The presentation could also be seen on the off-camera teleprompter allowing the teacher to focus their eyes on the target, avoiding the student feeling the teacher was only interested in the table device. This minimized the possibility of the student getting distracted.

Finally, the practical sessions were developed through screen-casting technology. More precisely, the teacher's computer screen is captured and the student can see all the actions made by the teacher while she simultaneously records her voice and explaining everything that is taking place. For this kind of session a good microphone is also essential.

---

[8] https://www.edx.org/blog/how-mooc-video-production-affects
[9] Using an application like *ExplainEverything*

Platform

As has already been mentioned, the course was deployed on the CANVAS[10] '[11] platform. Canvas is a new, open-source Learning Management System (LMS) that has several interesting properties aligned with our interests. Firstly, it is fully open-source making it suitable for installation on the university servers, allowing full control over course activity, as well as the ability to define new metrics for assessing the course. In addition, it has a strong social-network focus as it borrows a number of features from, for example, Facebook, Twitter, and Youtube (Godwin-Jones, 2012). Furthermore, it allows you to develop your own modules (3rd API tools) with Learning Tools Interoperability (LTI) standard or specific canvas APIs. In terms of development, it has a "dev & friends" site which is a great resource for getting started in developing on Canvas itself as it contains forums and documentation indexed by a search engine.

We implemented two different 3rd-party apps: i) a customized video streaming server and a ii) software grading tool based on DreaMT[12] and Jutge.org (Petit, 2012) in a similar way to that developed by (Staubitz, Renz, Willems, Jasper, & Meinel, 2014).

As mentioned, the Jutge.org platform is used to automatically evaluate programming assignments. Jutge.org is an online programming judge mainly used at UPC to assist students and instructors in massive regular courses in the Computer Science and Mathematics schools. It currently offers more than 1500 problems and has corrected more than one million submissions so far. Its average response time is six seconds per submission.

The architecture of Jutge.org offers a web front-end that exposes all its capacities for its users, and a back-end that handles client submissions. As this MOOC used the Canvas platform in order to deliver its content and administrate its courses, only the back-end of Jutge.org was required. In our case, the integration between this back-end and Canvas was solved using LTI[13].

Internally, Jutge.org's back-end is designed as a distributed master-slave system, with a master server that offers a queue into which clients post submissions and retrieve corrections, and several slave servers that host the different virtual machines that perform the actual correction tasks in a secure environment (Petit, 2012). In order to correct a submission, a virtual machine receives the problem description (including the solution written by the problem setter), the candidate solution and a driver module, which guides the compilation, execution, and checking process to produce a correction.

Note that executing arbitrary code submitted by arbitrary users is an important security concern. In order to isolate possible threads, see (Forisek, 2006) for a clear summary, the correction system of Jutge.org uses virtual machines with immutable drives that continually reboot and are placed behind a network firewall.

---

[10]http://www.instructure.com
[11]https://github.com/instructure/canvas-lms/wiki
[12]http://alopez.github.io/dreamt/
[13] http://www.imsglobal.org/toolsinteroperability2.cfm

Thanks to the modular architecture and the secure environment that Jutge.org offers, handling problems for this MOOC were solved easily; it was only necessary to create a new driver to guide the correction of MT problems. Since this driver shares many similar capacities to the already-existing ones in Jutge.org, developing it took only a couple of days.

Team and Promotion

As follows we explain the details on the course's team, the call for participation, and the motivational strategies.

Team. The course was developed in the context of the MOOCs 2013 contract with the AGAUR. The core team comprised people with PhDs with different profiles. The course was taught by Marta R. Costa-jussà (NLP) and Lluís Formiga (Speech Processing). Internal collaborators included Prof. José A. R. Fonollosa and Assoc. Prof. Francesc Vallverdú; our external collaborators were Meritxell González (MT Evaluation and Dialog Systems) and Jordi Petit (Software Engineering and Code Evaluation). In addition, Oriol Torrillas was one intern student involved part-time. Other very valuable support came from the Institut de Ciències de l'Educació (ICE) and MediaTic Terrassa.

Promotion. This MOOC was promoted using a low-resource strategy. We followed, to some extent, similar procedures to the announcement of Workshops and Conferences. Before any marketing took place, we set up a web page[14] to house the promotional video. Basically, this video showed a dialog between the MOOC lecturers and a selection of images presenting the university and the course content to motivate students to enroll. The dialog between the teachers was recorded under the same conditions as the video-lectures. In addition to the promotional video, the web page presented the structure of the course and the experiences of some students that had completed the course before its launch. Using the web page as the central reference point, the course was promoted via social media, including LinkedIn, Facebook, and Twitter, in addition to well-known distribution lists in the field, including corpora@uib.no, mt-list@eamt.org, and moses-support@mit.edu.. An announcement was placed on the NLPPeople web page for jobs in artificial intelligence, personal invitations were extended to universities and institutes, such as the Instituto Politécnico Nacional (Mexico), Universidade de São Paulo, and the Institute for Infocomm Research (Singapore), and the course was promoted in the Spanish University system, including paid leaflets at the SEPLN[15] conference. Additionally, a video teaser was available on YouTube[16].

Motivation. During the MOOC we developed strategies to keep students motivated and active. The main actions consisted of: (1) email announcements at the beginning of each week giving the content, at the deadlines for quizzes and assignments, and to enable participants to gain the certification for the course. The former and the second one followed MOOC's protocols

---

[14] Please, visit http://www.mt-mooc.upc.edu
[15] Sociedad Española para el Procesamiento del Lenguaje Natural, http://www.sepln.org
[16] https://www.youtube.com/watch?v=6XS1oP9gKV4

from large platforms such as Coursera; (2) revitalization of discussions; and (3) uploading of interesting and related links for each week's topics.

Student Details

Students in the course were from all over the world. Prior to week one,, we offered a pre-course survey that remained open throughout the duration of the course. A total of 348 students completed this survey. Questions were related to motivation, demographic details, background, and previous experience with MOOCs. The questions asked are reported in Table 2. Most of the questions were influenced by previous MOOCs taken by the professors of this course from popular platforms such as Coursera. The answers to these questions allowed the researchers to adapt the course to the MOOC audience. For example, questions about the background were really important to make on-the-fly, from which we could slightly adapt activities or modules, and give extra bibliography for example to really low programmer profiles.   In the following subsections we carefully describe the answers to these questions. Although we had a short reaction time between the collection of answers and the starting of the course, we explain in most of the cases how we took advantage of this valuable data.

Table 2

Pre-Course Survey Questions and Corresponding Figures

Topic	Question	Figure
Motivation	How did you hear about this course?	2
	Why are you taking this course?	3
Demographic	Year of Birth	4
	Gender	-
	Country of Birth	5
	In which country do you currently live?	6
Background	What kind of background do you have?	8
	What kind of programming background do you have?	9
Studies	What is your background in MT?	10
	Highest level of education completed?	7
	Which is your professional situation?	11
Work	In what type of organization are you?	12
	Teacher/lecturer?	-
	What type of courses have you participated in?	13
	How many free on-line courses have you participated in over the last two years?	14
	How many MOOCs have you participated in?	15
	How many MOOCs have you finished?	16

MOOC	In the cases where you finished the course, did you ask for a certificate?	-
	Have you ever paid to get a certificate?	-
	If you have not been able to finish a MOOC, how do you feel about that?	17
	What is the main motivational reason that led you to enroll in this kind of course?	18
	Which aspects convinced you to select a course?	19
	Was the content of the course related to your previous studies?	-
	In which platform(s) were the course(s) offered?	20
	In which language were the courses taught?	-
	Do you plan to take more MOOCs in the future?	-

Motivational Details

The main two motivational questions asked to the students were about how they found out about the course and why they were taking it. Results are shown in Figure 2 and 3, respectively. Around 50% of students found out about the course through a colleague or a friend. Around 50% of students did the course for educational reasons and 30% for working reasons. From the very beginning, the course was designed to cover both theoretical and practical issues. Therefore, from the motivational side, students were exactly as expected.
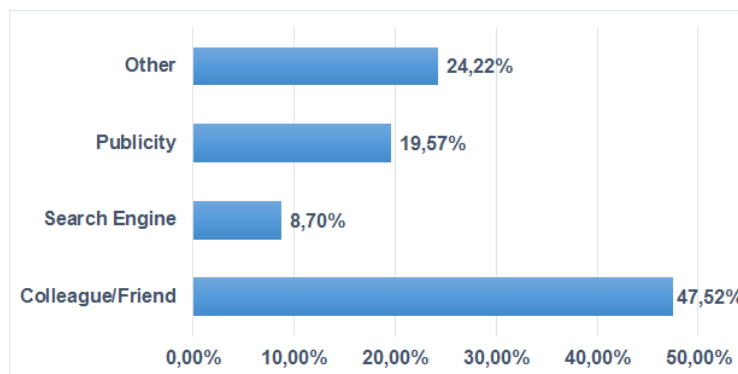
Figure 2. How did you hear about this course?

186
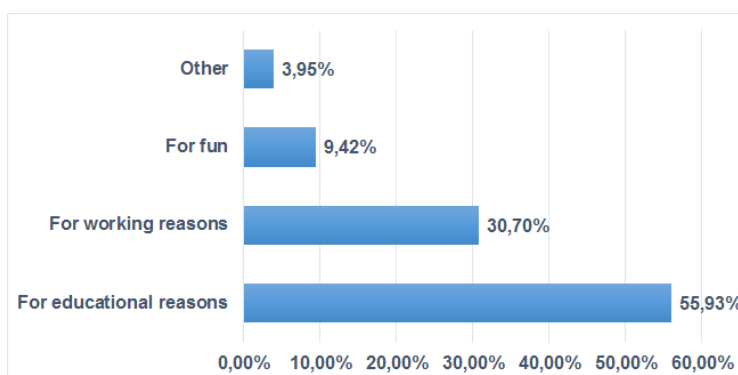
Figure 3. Why are you taking this course?

Demographic Details

Figure 4 shows the distribution of the students in relation to their year of birth. The variance is quite large, but most students were between 25 and 35 years old; 42% were women and 58% were men. Young people have higher flexibility in learning as supported by many different studies (Chang, Shibata, Andersen, Sasaki, & Watanabe, 2014), which is a nice advantage for our group of students.
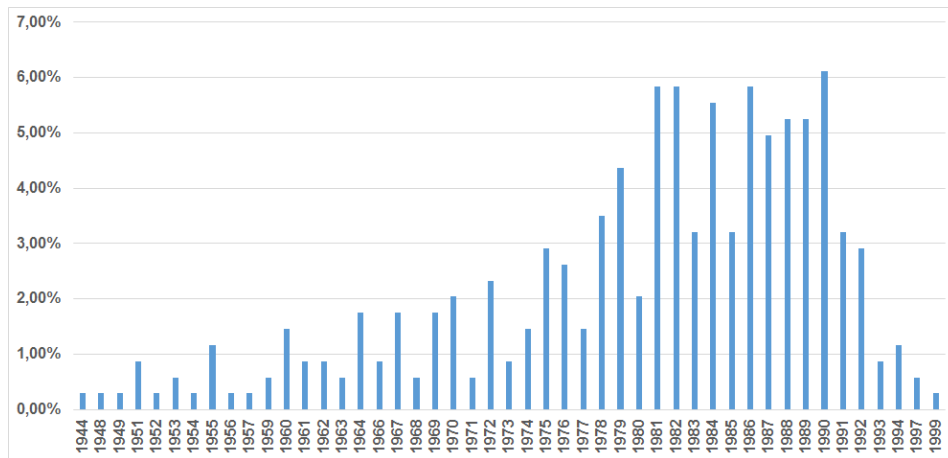
Figure 4. Year of Birth

Figure 5. Shows the distribution of the students in relation to their countries of birth. Figure 6 shows the distribution of the students according to the country they currently live in. Among most predominant countries we found Spain and Germany. Given that the course was promoted from a Spanish University, it makes sense that the course has a larger number of students from this country. However, we cannot give an explanation of the large number of German students. The group of other countries include: Algeria, Argentina, Austria, Azerbaijan, Bangladesh, Belarus, Belgium, Bosnia, Brunei, Bulgaria, Cambodia, Canada, China, Colombia, Croatia, Czech Republic, Caribeean, Egypt, Hungary, Indonesia, Iran, Iraq, Ireland, Israel, Japan, Jordan, Kazakhstan, Latvia, Morocco, Netherlands, Norway, Palestine, Paraguay, Perú, Philippines, Poland, Portugal, Romania, Arabia Saudi, Serbia, Singapore, Slovenia, Somalia, South Africa, South Korea, Srilanka, Switzerland, Tunisia, Turkey, Ukraine, Uzbekistan, Venezuela, and Vietnam.
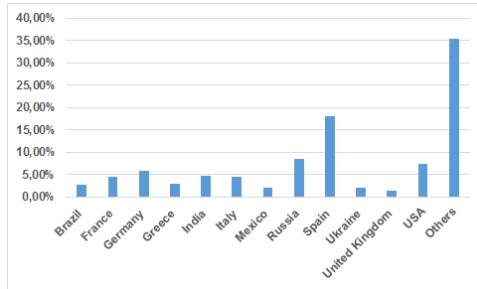
Figure 5. Country of Birth
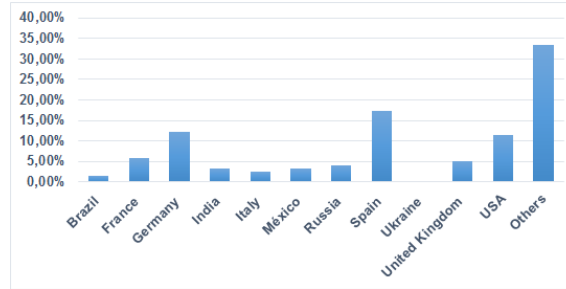
Figure 6. Country of Living

As for these demographic details, we were pleased to have such a variation on the country of living, because it gave a higher international dimension to the course. We can consider that we did a nice promotion given the coverage of most continents. In further editions, we have to make sure that we cover Australia properly as well.

Background

We asked several background questions to better comprehend what students expected from the course.

Studies. Figure 7 shows the distribution of the students according to the highest level of education completed. Figure 8 shows the range of the students' backgrounds. Around 50% of students have a master degree, achieving the desired student profile. The kind of backgrounds is mostly between linguists and computer scientists, which was what we expected when designing the course. The variability between these two profiles points out a challenging issue in our course. We have to adapt to people with different programming skills, as we show in the following questions.
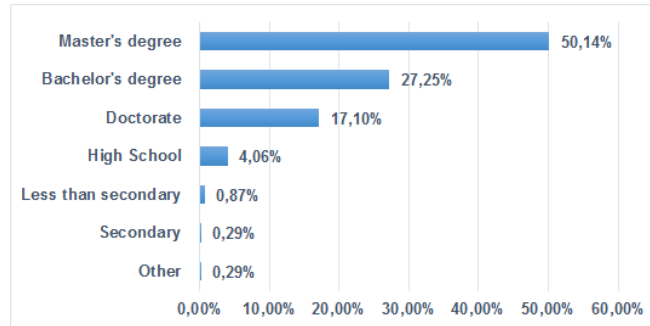
Figure 7. Highest level of education
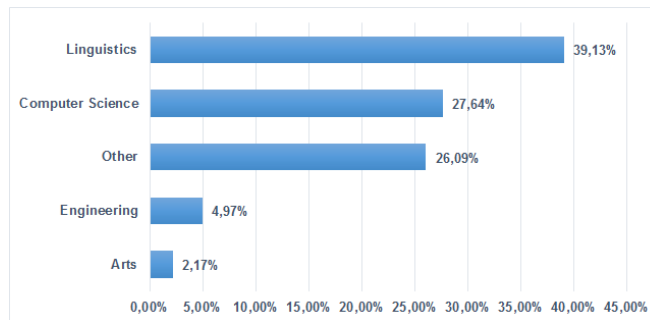
Figure 8. Kind of background.

Figures 9 and 10 show the students' background in programming and MT, respectively. Most students were programming beginners and without MT experience. This allowed us to expect more questions in the programming assignments and put more effort on explaining further details on each of them.
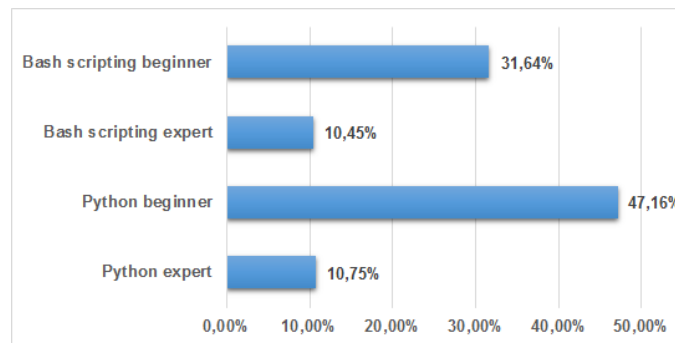
Figure 9. Programming background.
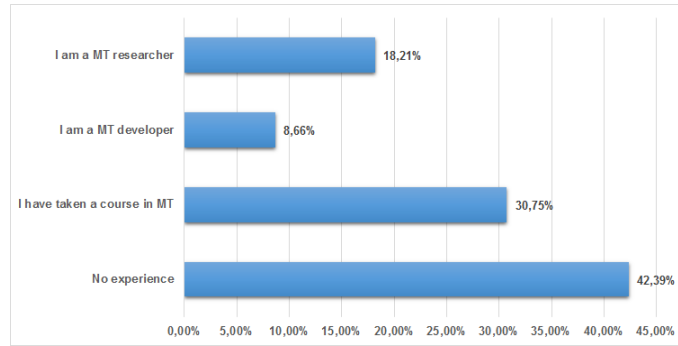
Figure 10. MT background

Work. Figure 11 shows the distribution of the students related to their professional situation. Figure 12 shows the distribution of the students according to the type of organization they work in. Student professional situations are shared between private and public sectors, while the type of organization students affiliate with vary between educational and industrial. Again, this also motivated us towards balancing theory, practice, and applications during the entire course.
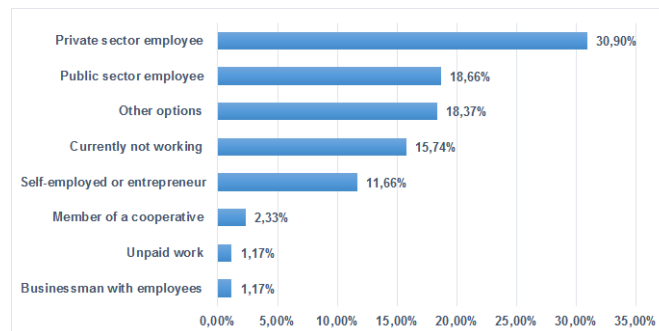
Figure 11. Professional situation
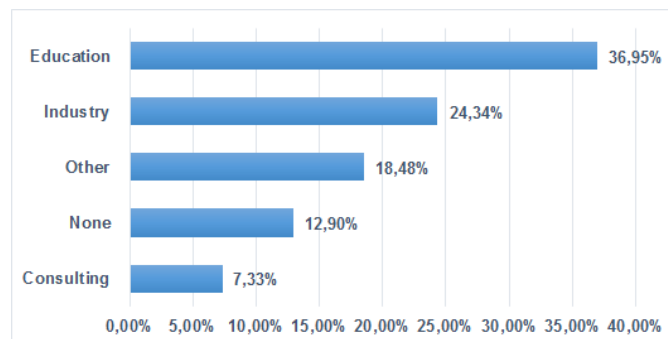
Figure 12. Type of organization.

Finally, the percentage of teachers/lecturers was 22% compared to 78% non-teaching participants. This is not a course where we expected a large number of teachers/lecturers given its nature. However, we had quite a high percentage.

Experience with MOOCs

Figure 13 shows the distribution of the students in relation to the type of courses they have participated in. Figure 14 shows the distribution of the students according to the number of free on-line courses they have participated in over the last two years. Figure 15 shows the distribution of the students in relation to the number of MOOC courses they have participated in. Figure 16 shows the distribution of the students according to the number of MOOC courses they have finished. Surprisingly, this was the first MOOC for a lot of the students. This last finding is kind of positive and negative at the same time. On the one hand, students would not compare to other courses. On the other hand, students would not be used to the procedures of on-line courses.
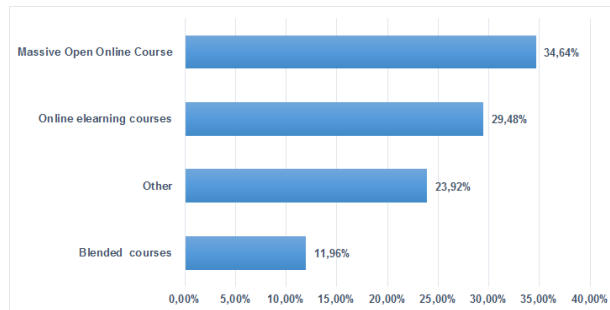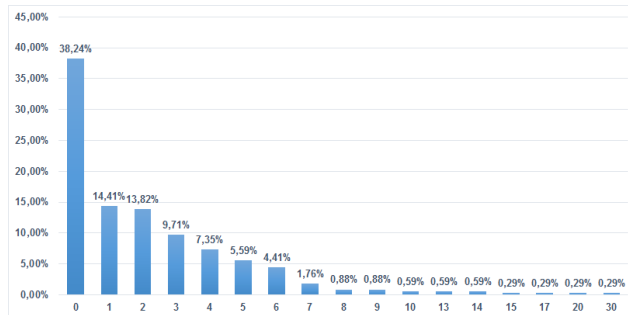

Figure 13. Type of courses been involved.
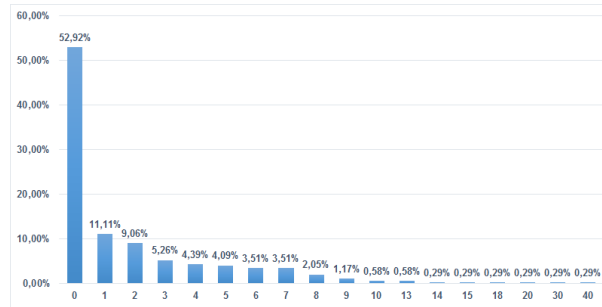

Figure 14. Free on-line courses taken.

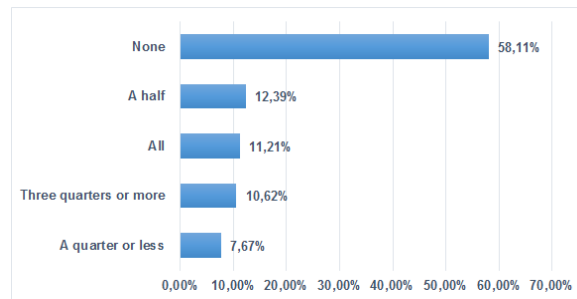*Figure 15.* Number of MOOC courses that students have participated.

Figure 16. Number of MOOC courses that students have finished

In the 44% of cases, students did not finish the courses, and in 82% of cases, students have never paid to get a certificate. Statistics of not finishing the course does not favor us. However, the high number of students who did not pay for a certificate supported our decision for a 100% open and free course, which was advertised from the very beginning on the promotional video. Figure 17 shows the distribution of students in relation to how they would have felt if they had not been able to finish the MOOC. Figure 18 shows the distribution of the students according to how the students are motivated to enroll in this kind of course.
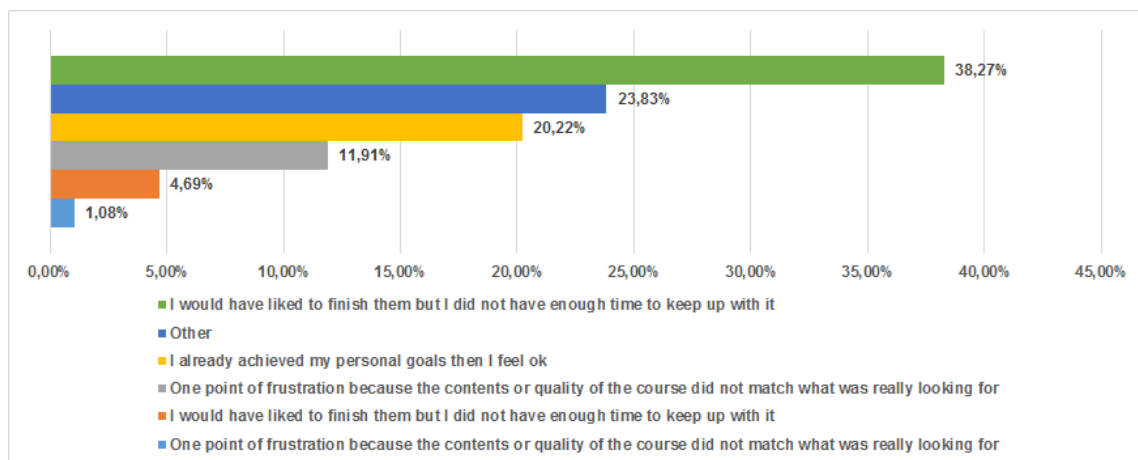

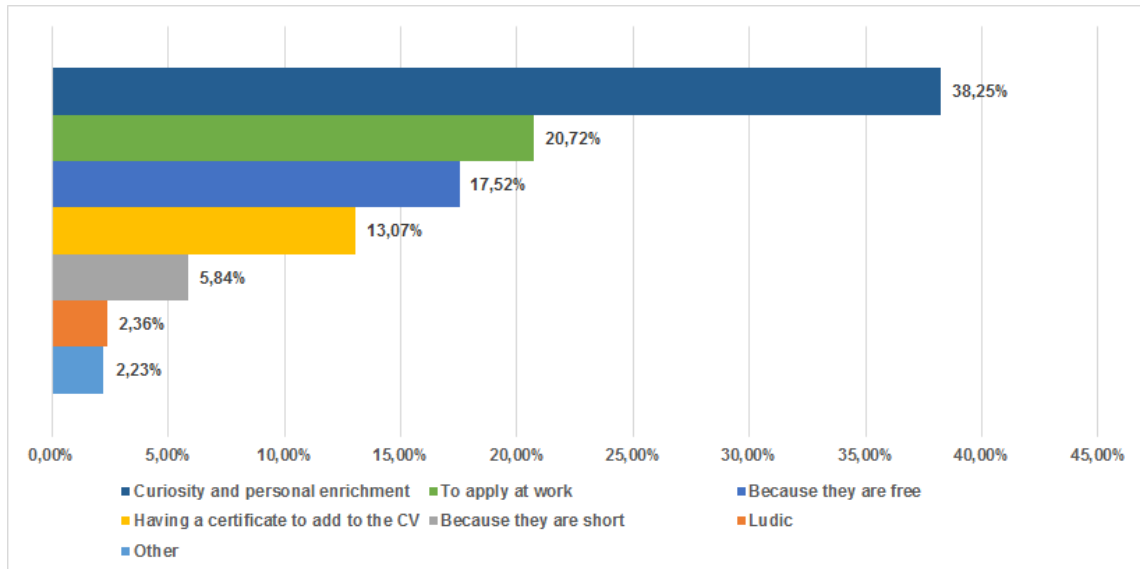*Figure 17.* Feelings about not being able to finish the MOOCs.

192

Figure 18. Motivation to enroll in this kind of courses

Figure 19 shows the distribution of the students in relation to the aspects that convinced them to select a certain course. In almost 75% of cases the students stated that the courses they chose was related to their previous studies. Figure 20 shows the distribution of the students according to the MOOC platforms they have used, Coursera being the most prevalent in this case.
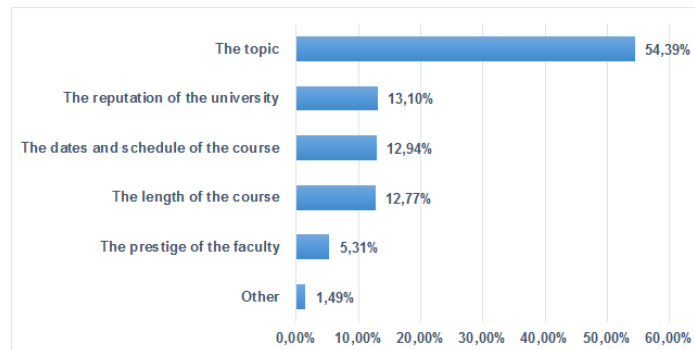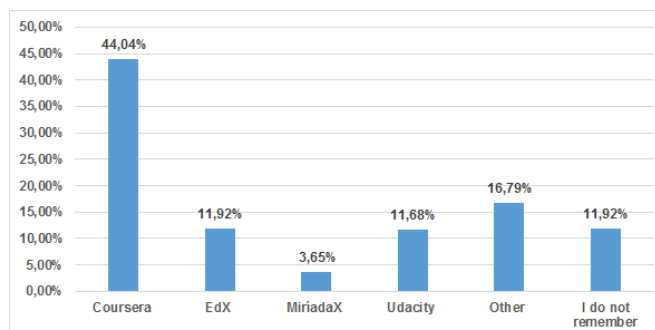

Figure 19. Aspects to select a course.

Figure 20. Experience with MOOC platforms.

In 77% of cases, the MOOC courses were in English, and in 1% of the cases, the MOOC courses were in Spanish. Although, large platforms such as Coursera are offering more and more courses in languages other than English, English continues to be the most widely used. In our case, given the nature of our subject, which was quite advanced, we considered it adequate to use English which is the language most used for this type of research. Little by little, new technologies such as Speech Recognition and, precisely, MT, will make automatic subtitles easier to implement for MOOCs[17]. Finally, 61% of the students stated that they would take further MOOCs in the near future.

Participation, Feedback and Correlations

In this section we report the participation of students during the MOOC, as well as the feedback received by means of a post-course survey. Finally, we analyze several characteristics of successful students and show the corresponding correlations with the entire group of students.

Participation

As expected, participation was quite high during the first week of the MOOC and started to decrease as the weeks progressed. This is a fairly standard pattern in MOOCs that we tried to prevent by using the motivational actions mentioned in Section 6. Table 3 shows the statistics of video-lecture views. Note that views decreased from the first week onwards; however, this stabilized over the final fortnight. The average views per week are reflected in Figure 21.

---

[17] https://support.google.com/youtube/answer/3038280?hl=en

Table 3

Statistics of Video Lectures Including Week Number, Duration in Minutes and Number of Views.
Video-Lectures in Italics were Optional.

W	Topic	Min	Views
0	Forthcoming MOOC on Approaches to Machine Translation	2:50	1,342
1	Introduction: Motivation, history and applications of MT	8:30	627
	Introduction: MT Approaches	11:40	535
	Introduction: SMT Challenges	8:02	404
	Rule-based MT	11:46	424
2	Alignment (I): Introduction	13:44	449
	Alignment (II): IBM Model 1 and EM	15:34	333
	Alignment (III): IBM Model 2 and HMM	12:49	265
	Alignment (IV): IBM Higher Models	17:26	127
	Alignment (V): Applications, Quality and Problems of Word Alignment	11:08	108
	Growing Heuristics, Symmetrization and Phrase-Models	12:33	190
	Phrase-Extraction and Scoring	19:55	176
3	Reordering	4:55	139
	Language Modeling (I): Introduction	17:34	118
	Language Modeling (II): Modeling	11:36	107
	Language Modeling (III): Modeling	12:00	88
	Language Modeling (IV): Applications to MT	9:04	86
	Tuning parameters	16:02	80
	Domain Adaptation	7:56	78
	Decoding (I)	5:24	90
	Decoding (II)	4:57	95
4	Evaluation (I): Intro	9:07	101
	Evaluation (II): Automatic Metrics	12:53	81
	Evaluation (III): Human Metrics	6:16	64
	Evaluation (IV): Advanced Metrics	11:52	70
	Asiya: an Open Framework for MT Evaluation and Output Analysis	11:35	85
	Asiya Online + Demo	6:35	77
	Asiya tSearch + Demo	10:01	98
	Toolkits	9:27	53
	Asiya: Web Service	2:35	53
5	Hybrid MT: Orthography and Lexis	7:00	76
	Hybrid MT: Morphology (I)	10:17	62
	Hybrid MT: Morphology (II)	9:20	59
	Hybrid MT: Syntax (I)	6:30	50
	Hybrid MT: Syntax (II)	7:23	51
	Hybrid MT: Syntax (III)	6:37	54

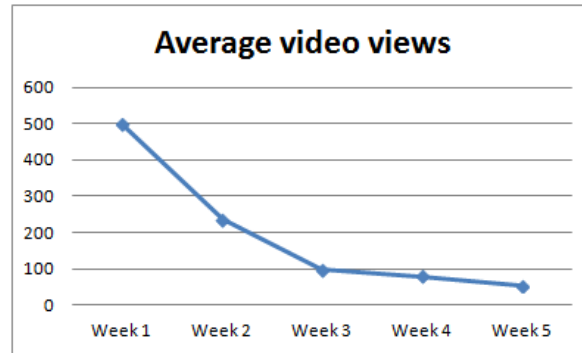|		
Hybrid MT: Semantics (I)	5:30	49
Hybrid MT: Semantics (II)	10:50	47
Case study: Computer Aided Translation	24:23	45
Conclusions	1:59	44

Figure 21. Average video views per week.

Table 4 shows the statistics of quizzes, assignments, and the final exam including participation, minimum, maximum, and mean scores.

Table 4

Statistics of Participation for Quizzes, Assignments and Final Exam Including Minimum, Maximum and Average Scores.

Activity	Participant	Min	Max Score	Average
Quiz 1	235	0	5	3.6
Quiz 2	91	0	6	4.3
Quiz 3	78	0	4	3.0
Quiz 4	49	1.6	8	6.6
Quiz 5	48	1	5	4.1
APA	31	0.0	6.1	4.0
EPA-LM	34	0.0	8.5	3.0
EPA-Asiya	42	2	18	13.5
Final	40	0	15	9.5

The discussion folder received over the entire course included a total of 87 threads and 287 replies. The average activity time per student was 16 hours 29 minutes, without counting the 211 students that never began the course or entered the forums.

The course had a total of 34 successful candidates that received the certificate, meaning the course achieved an almost 5% completion rate counting all students registered and a more than 6% completion rate when subtracting from the total number of students those that never began the course (211). This result confirms average MOOC completion rates of less than 7% [18].

Feedback

As follows, this subsection reviews the responses received in the post-course survey. The survey opened five days before the end of the course to find out students' opinions of the quality of the course, the video-lectures and assignments. Table 5 reports the questions posed in this survey. A total of 42 students took this survey, which is coherent with students that were actively following the course at this stage.

Figure 22 shows whether students found the course useful. As we can see from the responses, most students enjoyed the course (over 88% where on the positive and neutral side). However, we cannot ignore the fact that 12% of the students disagreed or strongly disagreed on finding the course useful. Looking at the profile of those students, we discovered that they were the students with higher backgrounds in research, programming, and MT. In the future, to reduce this percentage, we have to further clarify the level of MT that the course aims to achieve. Figure 23 shows what areas of MT students were most interested in. The highest level of interest was divided between hybrid and statistical MT. However, there was a significant percentage of students interested in rule-based and evaluation. There was, therefore, no topic that could be taken out of the syllabus.

Table 5

Post-Course Survey Questions and Corresponding Figures

Question	Figure
Did you find this course useful?	22
In which type of MT were you most interested?	23
I enjoyed being a part of this course.	-
What I really missed in the course is…	-
Approximately how many hours per week did you spend on the course?	24
Approximately how many hours per week did you spend in the forum?	25
Video-lectures were carefully prepared and well-deployed	26
Assignments were carefully designed, explained and adequate for the level of the course	27

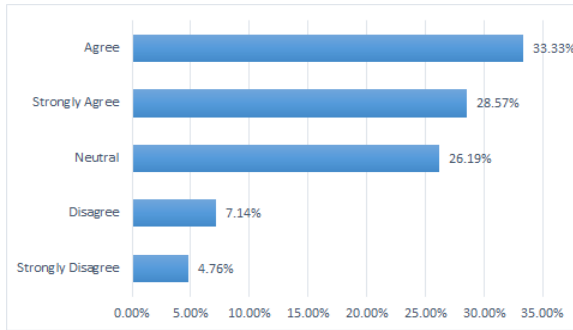[18] https://www.insidehighered.com/news/2013/05/10/new-study-low-mooc-completion-rates

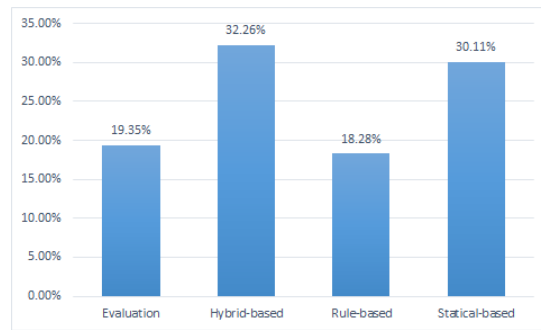*Figure 22*. Did you find this course useful?

Figure 23. More interested type of MT.

More than 72% of the students enjoyed taking part in the course. Figures 24 and 25 show the time per week that the students spent on the course and in the forum, respectively. Here, we see that the majority of students spent between 2-4 hours per week doing the course. However, there was a very significant percentage of students (35%) that had to dedicate more than 4 hours to the course. This could be due to the different backgrounds of the course participants and the difficulties of adapting study hours to such a variety of profiles. On the other hand, time spent in the forum was quite low.

Figures 26 and 27 show the students' opinions of the quality of the video-lectures and assignments, respectively. Most students agreed the course offered good quality video-lectures and assignments. However, while the feedback from the students was positive, it should be noted that the 67% of the feedback was received from successful students.
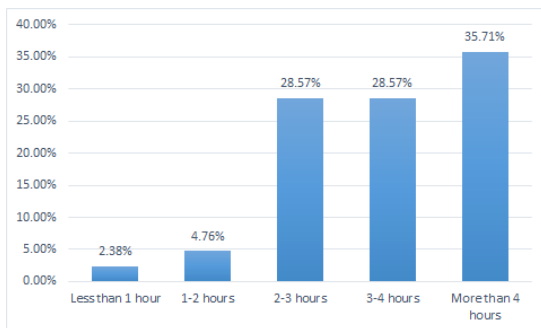

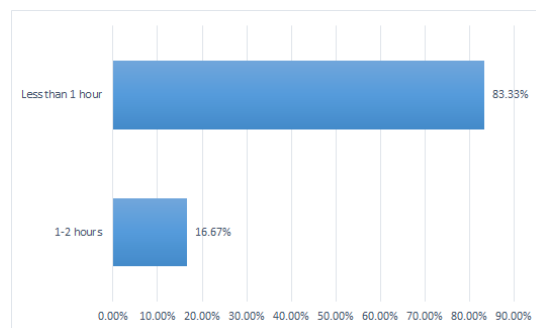*Figure 24*. Hours per week in the course.

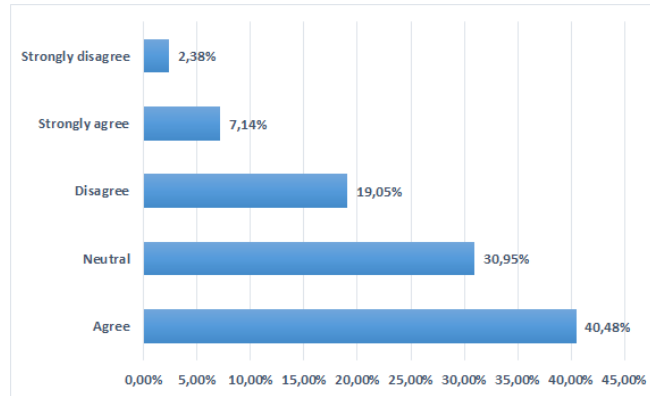*Figure 25*. Hours per week in the forum.

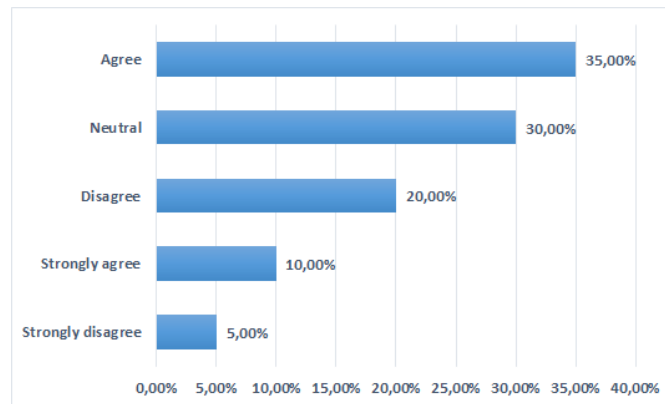*Figure 26*. Video lectures were carefully prepared.

Figure 27. Assignments were carefully designed.

Correlations

We have previously described the global statistics relating to the students who took the course. In this section we analyze the successful students in correlation with the total number of students in the most relevant demographic groups.

For the successful students, the motivation for doing the course tended to lean towards educational reasons (70%) as compared with the total number of students that were more inclined to take the course for work-related reasons.

Figure 28 shows the distribution of the students in relation to their year of birth. In green, we show the percentage of successful students per year of birth. There were more successful students in the younger profiles. 48% of the students were women and 52% were men. Slightly more female candidates successfully completed the course than their male counterparts.

Figure 29 shows the distribution of the students in relation to their country of birth. Figure 30 shows the distribution of the students according to the country they currently live in. In both figures, green shows the percentage of successful candidates per country of birth/living. Mexico had the biggest percentage of successful students.
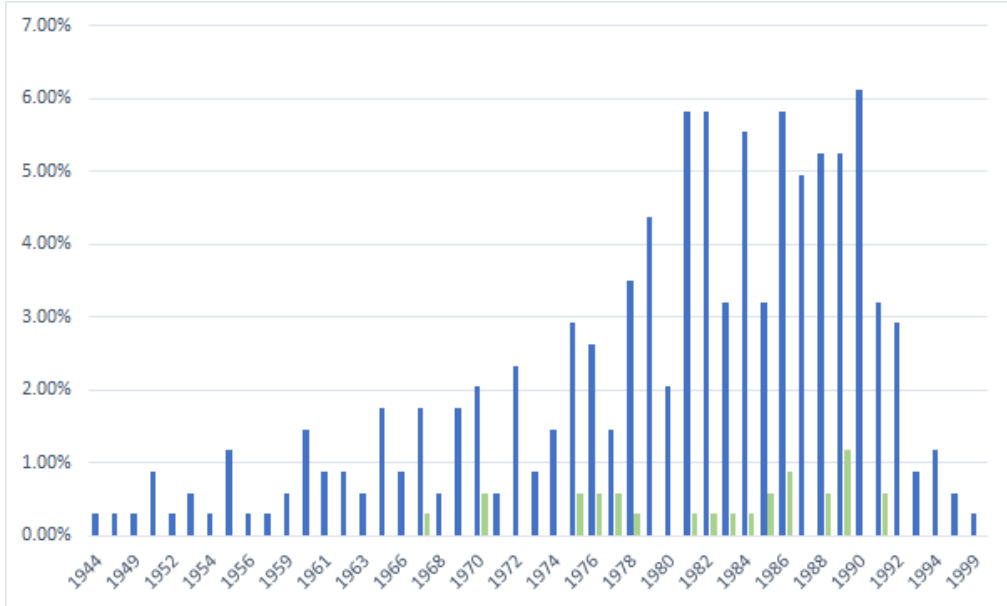

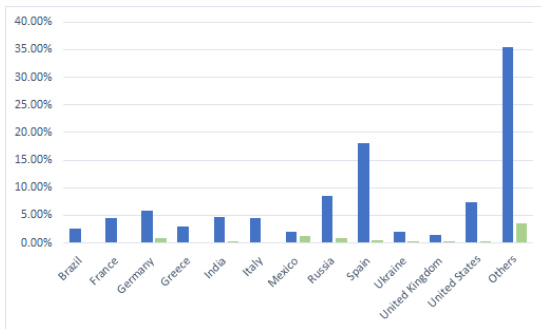*Figure 28.* Year of Birth of successful students: in green, percentage of successful students.

Figure 29. Country of Birth of successful students: In green, percentage of successful students.
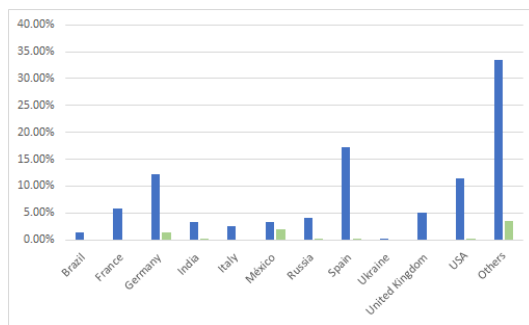
Figure 30. Country of Living: In green, percentage of successful students.

Figure 31 shows the distribution of the students in relation to the highest level of education achieved. In this case the distribution is very similar to the case of total number of students. Figure 32 shows the range of the students' backgrounds. In this case, the profile of students changes with respect to the profiles of the total number of students. Most successful students had

a background in computer science rather than linguistics, which may support some of the discussions in the forum around the fact that people with linguistic profiles thought that this course was more oriented to computer scientists. The topic of the course may be weighted towards computer science, but, we intended to open it up to different profiles, which resulted in a 54% success rate with students not holding a computer science background.
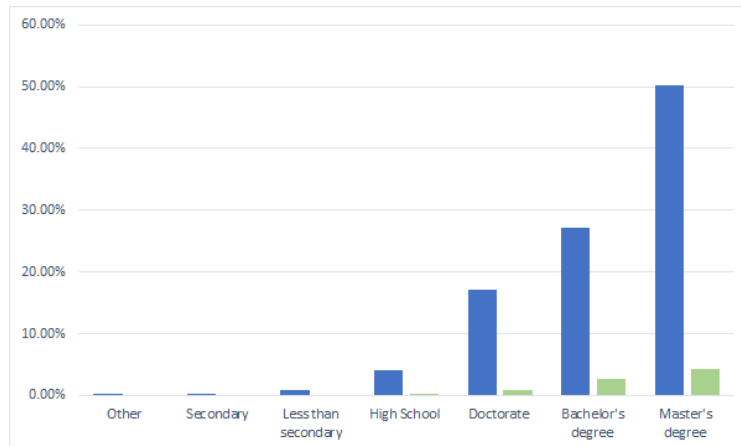

Figure 31. Highest Level of Education of successful students: in green the percentage of successful students.
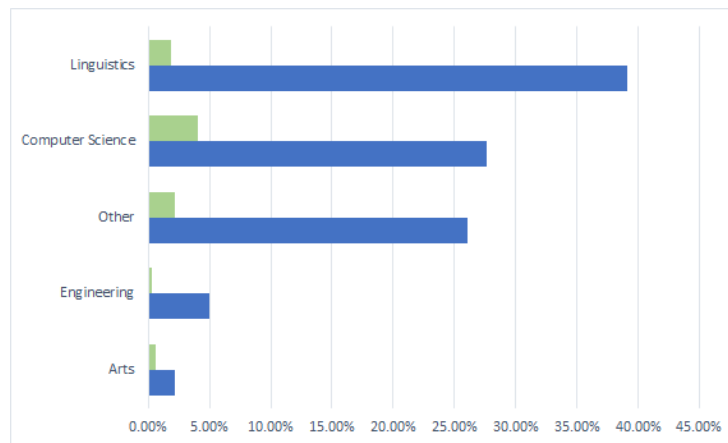

Figure 32. Kind of background of successful students: in green percentage of successful students.

The correlation between successful students and a background in programming and MT is almost 100%. Only 18% of the students were lecturers and equal proportions, 35%, came from the education and industry fields. Seventy percent of participants confirmed that this was their first MOOC, and about 56% of students confirmed that they would take more MOOCs in the future.

Summary and Conclusions

This paper reports the details of the first Massive Open On-line Course held at the Universitat Politècnica de Catalunya. The efforts involved in setting up this first course of its kind were made in many areas including the establishment of a new platform, international promotion, the recording of videos, and the definition of activities. Most of the design has been influenced by previous works in the MOOC area as cited along the paper. The description, procedures, and analysis of the MOOC reported in this paper can be especially useful to professors and/or researchers that want to prepare a MOOC on a low-resource basis and a less standardized platform. To this end, we show a summary table of what did work, what more or less worked in our experience, and the future work.

Table 6.

Summary of What Did Work and What Did Not Work in the MOOC Design.

	Worked	**More/Less Worked**	**Future Work**
Content	Mix Theory and Practice State-of-the-art Research	Content addressed to different backgrounds -- Less mathematical | More empirical	Provide perspective from Translation companies and professionals
Activities and Forum	Quizzes On-line autoevaluated coding	Programming Activities Peer2Peer cooperation	Better social network perspective
Recording Set	Format Slide + Professor	Applications for writing in the slides vs wacom	Student-given seminars Hangouts
Platform	User-friendly	LTI interconnection	Deal with other platforms
Team	Multidisciplinary		
Promotion	Distribution lists Specialized Web Pages Personal WebPage	Paid advertisements Geographical impact of marketing	More wide promotion in companies environment

Pre-Course Survey	Design	Not enough time to reach to the type of audience	Make it mandatory for evaluation
Post-Course Survey	Design	Not enough participants	Make it mandatory for evaluation Do a middle-course survey
Motivation Strategy	Weekly communication with students	Generic emails to remind lectures, activities and evaluation Failure prediction	Keeping personalized emails with progress track. Student-based prediction models

The course, which is on the very specific topic of Machine Translation, was well received by the international community, and students from all over the world took part. Details on their participation were gathered and analyzed throughout the pre-course and demographic surveys. The completion rate falls within the standard for this kind of course (below 7%), but the post-course surveys gave very positive feedback. In the future, we plan to do further editions of this course giving special focus on achieving a higher success rate of completion by selecting motivated students and making further actions to follow their learning.

Acknowledgments

References

Abeer, W., & Miri, B. (2014). Students' preferences and views about learning in a MOOC. *Procedia - Social and Behavioral Sciences*, *152*, 318–323. doi:10.1016/j.sbspro.2014.09.203

Aydemir, M., Özkeskin, E. E., & Akkurt, A. A. (2015). A theoretical framework on open and distance learning. *Procedia - Social and Behavioral Sciences*, *174*(2003), 1750–1757. doi:10.1016/j.sbspro.2015.01.833

Chang, L.-H., Shibata, K., Andersen, G. J., Sasaki, Y., & Watanabe, T. (2014). Age-related declines of stability in visual perceptual learning. *Current Biology*, *24*(24), 2926–2929.

Costa-jussà, M. R. (2013). Morphological , syntactical and semantic knowledge in statistical machine translation. In *NAACL HLT 2013 Tutorial Abstracts* (pp. 16–18). Atlanta, USA: ACL.

Costa-jussà, M. R., & Farrús, M. (2014). Statistical machine translation enhancements through linguistic levels. *ACM Computing Surveys*, *46*, 1–28. doi:10.1145/2518130

Costa-jussà M., Formiga, L., Petit, J. & Fonollosa, J.A.R. (2014). Detailed description of the development of a MOOC in the topic of Statistical Machine Translation *Proc. of the Mexican Conference on Artificial Intelligence, MICAI*, Chiapas.

Gil, R., Virgili-Gomá, J., García, R., & Mason, C. (2015). Emotions ontology for collaborative modelling and learning of emotional responses. *Computers in Human Behavior*, 0–7. doi:10.1016/j.chb.2014.11.100

Giménez, J., & Màrquez, L. (2010). Asiya: An open toolkit for automatic machine translation (meta-) evaluation. *The Prague Bulletin of Mathematical Linguistics*, *94*, 77–86. doi:10.2478/v10108-010-0022-6.PBML

Godwin-jones, R. (2012). Challenging hegemonies in online learning. *Language Learning & Technology*, *16*(2), 4–13.

Harms, R. (2015). Self-regulated learning, team learning and project performance in entrepreneurship education: Learning in a lean startup environment. *Technological Forecasting and Social Change*. doi:10.1016/j.techfore.2015.02.007

Koehn, P. (2010). *Statistical Machine Translation*. Cambridge University Press.

Kolp, M., & Wautelet, Y. (2014). Computers in human behavior human organizational patterns applied to collaborative learning software systems. *Computers in Human Behavior*, 1–10. doi:10.1016/j.chb.2014.11.094

Lopez, A., Post, M., Callison-burch, C., Weese, J., Ganitkevitch, J., Ahmidi, N., Buzek, O., Hanson, L., Jamil, B., Lee, M., Lin, Y., Pao, H., Rivera, F., Shahriyari, L., Sinha, D., Teichert, A., Wampler, S., Weinberger, M., Xu, D., Yang, L. and Zhao, S. (2013). Learning to translate with products of novices : A suite of open-ended challenge problems for teaching MT, *Transactions of the Association for Computational Linguistics 1*, 165–178.

Muñoz-Merino, P. J., Ruipérez-Valiente, J. a., Alario-Hoyos, C., Pérez-Sanagustín, M., & Delgado Kloos, C. (2015). Precise effectiveness strategy for analyzing the effectiveness of students with educational resources and activities in MOOCs. *Computers in Human Behavior*, *47*, 108–118. doi:10.1016/j.chb.2014.10.003

Petit, J., Giménez, O., & Roura, S. (2012). Jutge.org: An educational programming judge. In *Proceedings of the 43rd ACM technical symposium on Computer Science Education - SIGCSE '12* (p. 445). doi:10.1145/2157136.2157267

Rendas, Pinto, P. R. and Gamboa (1999), A computer simulation designed for problem-based learning. *Medical Education*, *33*: 47–54. doi: 10.1046/j.1365-2923.1999.00261.x

Savin-Baden, M. (2000). P*roblem-based learning in higher education: Untold stories.* UK: McGraw-Hill Education.

Savin-Baden, M., & Wilkie, K. (2006). Problem-based learning online. UK: McGraw-Hill Education.

Shen, C., & Kuo, C.-J. (2015). Learning in massive open online courses: Evidence from social media mining. *Computers in Human Behavior*. 1, 568–577 doi:10.1016/j.chb.2015.02.066

Staubitz, T., Renz, J., Willems, C., Jasper, J., & Meinel, C. (2014). Lightweight ad hoc assessment of practical programming skills at scale. In *IEEE Global Engineering Education Conference, EDUCON* (pp. 475–483). doi:10.1109/EDUCON.2014.6826135

Yang, Z., Liu, Z., Liu, S., Min, L., & Meng, W. (2014). Adaptive multi-view selection for semi-supervised emotion recognition of posts in online student community. *Neurocomputing*, *144*, 138–150. doi:10.1016/j.neucom.2014.05.055