

Subtitles on everything for everyone

Enabling community subtitling and translation for every video on the net.

1. Overview:

The potential: community-driven subtitling could make video global

Inside a community of interest, collaborative subtitling is a proven model. An individual or small group uses available tools to create a time coded transcript in the original language, which others can then translate. The first step can be quite time consuming, but the second can happen quickly. Viewers can search sites that serve their community for subtitle files in a language they speak, and use a compatible video player (like VLC) to display subtitles alongside video. The tasks, roles, and tools are well understood, and the process is applicable to anything from films to daily programming to short video clips. Video is already much better than audio or text at crossing global divides, because many popular videos are language-neutral. Community subtitling can cover the rest, making online video a truly global medium.

The problem: fragmented tools create fragmented communities

The problem with community subtitling is that the experience for viewers and contributors is highly fragmented; only the most motivated users overcome these barriers. The viewer's video player does not tell her anything about the availability of subtitles. Viewers must search the web, already knowing what they are looking for. When the viewer finds matching subtitles, she needs to pair them the video using a compatible video player. If the video is stored on the web in a proprietary player like Flash, this can be difficult or impossible.

Similar barriers face volunteer translators: when viewing a video, it is never clear where to find time-coded transcripts, or where to submit a new translation or correction. Different subtitle sites exist for different types of content and different languages. The end result of these hurdles is that subtitle communities only tackle the most sought-after videos, and even then, only serve the most enthusiastic and tech-capable viewers. New subtitling services are cropping up, but each exists as its own separate world: separate from the primary viewing context (e.g. Youtube) and separate from each other.

The solution: build open translation communities around the viewing experience

By connecting subtitle search and submission directly to the viewing experience, we can create a subtitling community around nearly every video online. By maintaining an open database of videos and matching subtitles, we can maximize the benefit of users' efforts. When the viewer watches a video, an indicator will tell her when subtitles are available, and in what languages. Once selected, the subtitles will display immediately, just like on a DVD. While watching a video, the user is one click away from offering feedback and corrections or submitting a new translation. Even better, we can create this experience around the vast majority of videos online. Every video podcast, every downloaded video, and every video on major video sharing sites would be covered, first in the open-source aggregator Miro (phase 1) and then in the Firefox browser

using a browser plugin (phase 2). Each piece can be open source and the pieces can communicate across open APIs; the system itself will be decentralized and open to collaboration.

Outcomes: more transcription, more languages served, better availability

We're confident that by connecting collaborative subtitle creation with the viewing experience we can create orders of magnitude improvements on every relevant metric. First, any subtitle that is created will become far more available once viewers can link those subtitles to a video with a single click. Second, translation is much easier than transcription, it can be divided among multiple translators or even wiki'd, and once a translation exists in a common second language like English, the pool of translators widens. If we make subtitles available to more viewers and give translators good ways to collaborate, the pool of translators will increase proportionally. Finally, the increased exposure and demand for subtitles will encourage more transcription, both by communities and by video creators themselves. Creating time-coded transcriptions is the most specialized and time-consuming part of the process. However, tools and collaboration methods are improving. We believe an increase in audience will bring more energy into the process, stimulate further improvement of collaborative tools for transcribing, and encourage creators of the original video to make time coded transcripts available. Massive amounts of energy goes into the creation, criticism, and promotion of online video. The tools we propose will bring that energy into the collaborative subtitling process and make collaborative subtitles a ubiquitous feature of the online video space.

Broad impact: the end of language barriers for popular video

The system we're proposing is the first step in an open, non-hierarchical structure that has the power to make any video available in any language where it has a motivated audience. Any video with a sufficiently large or particularly passionate audience will find volunteers ready to create a time-coded transcript and offer translations. This popularity threshold is unknown and will vary, but we do know that as translation communities grow and tools improve, the threshold will become progressively lower. Online video already exists in highly international, language-neutral spaces (because much popular content is language neutral). By giving viewers easy ways to find subtitles and participate in their creation, we can make language barriers porous.

2. How it works

The system consists of two parts: a subtitle fetcher built into a video player and a web-service. The web-service is a website where users maintain a database of videos and corresponding subtitles. The subtitle fetcher is a plugin (or feature) in a video player that checks this web service for available subtitles as the user watches video.

Subtitle fetcher in Miro (mockup). An indication appears below the video area when subtitles are available. When subtitles are available, the user can choose from a list of languages. Once selected, subtitles are downloaded and displayed in the video window. The user can visit the subtitle repository to search for subtitles and submit one, even if no subtitles are available. The "Find more" and "Submit new" links open a browser window and take the user to a subtitle repository website. Users can adjust how subtitles and notifications are displayed, and can choose to load subtitles automatically for preferred languages. The system will allow users to add multiple subtitle repositories, enabling a competitive ecosystem for subtitles with no centralized point of control (these repositories can, in turn, elect to share data with each other).



Subtitle fetcher Firefox addon on Youtube (mockup): The user installs a Firefox addon which uses a modified, subtitle-capable Youtube player to display each video. The Firefox addon will support Youtube first, followed by other popular video sharing sites. Sites can also choose to offer collaborative subtitling support directly to their users without requiring a Firefox addon (the addon is simply for ensuring compatibility with large sites like Youtube that might not offer such a service). The interface for choosing subtitles and participating is the same as in the Miro subtitle fetcher.

The screenshot shows a YouTube video player for the video "Obama Gives A Great Speech In Berlin, Germany". The video player includes a progress bar, a volume control, and a subtitle selector menu. The subtitle selector menu is open, showing the following options: English, Română, Svenska, Português, and Türkçe. Below the subtitle selector, there are links for "Find more" and "Submit new Options". The video has a "Rate" of 4 stars (390 ratings) and "Views: 27,756". Below the video player, there are links for "Share", "Favorite", "Playlists", and "Flag". There are also social media links for "MySpace", "Facebook", and "Digg". To the right of the video player, there is an advertisement for the Samsung Instinct phone, which says "NOW PLAYING. ONLY ON THE NOW NETWORK." and "SEE IT IN ACTION". Below the advertisement, there is a channel card for "The Young Turks", which says "From: TheYoungTurks", "Joined: 2 years ago", and "Videos: 1,948". There is a "Subscribe" button next to the channel card. At the bottom of the page, there is a "Commentary" section and a "Statistics & Data" section. The URL in the address bar is "http://www.youtube.com/watch?v=vVedDZU40Fo#".

Repository Site: The subtitle fetcher searches a subtitle repository website for subtitles matching a given video. The fetcher searches based on a range of available information (including hash, URL, title, and duration). Sometimes, as in the case of a matching hash or URL, the search will find an exact match (that is, we will be almost 100% sure that the subtitles indeed match the video). In this case, the fetcher will automatically acquire all available subtitles from the site and display these subtitles in the language selector shown in the mockups above.

However, even if a search by hash or URL comes up negative, it's possible that matching subtitles still exist (the video may have been re-encoded or reposted to another site, for example). In this case, users can click "Find more" in the subtitle fetcher, and be taken to the repository site to see a list of possible matches. If the user finds a match, they can create a new correspondence between the identifiers in their search (e.g. hash or URL) and the existing subtitle file. When they return to watch the video, the subtitle fetcher will load these new subtitles. In this way, the community

will build a list of links between subtitles for a given video and instances of that video across various sites and formats as the services progresses.

Clicking "submit new" in the subtitle fetcher will take users to a submission page on the repository site. If no subtitles for the video exist already in the system, the subtitle fetcher will automatically submit all available information on the video (title, hash, url, duration, etc) and users will be able to submit a new subtitle file. The submission page will feature links to tutorials on how to create time-coded transcripts using free tools. If subtitles exist already, the user will be able to submit a new translation. We will start with some simple capacity for moderating and improving translations (such as a wiki) and add more functionality for collaboration as we go.

3. Execution

(The order of this work might change depending on interest from the open source community or collaborating organizations / projects)

Phase 1: Subtitle fetcher in Miro, open APIs, open source alpha of repository site.

In the first phase we will create a working implementation of the system for all Miro viewers. Development of the subtitle fetcher in Miro is more straightforward (as compared to developing a Flash player / Firefox addon) so it makes sense as a first step. We will define the Open APIs through which the fetcher talks to the repository, and we will create and host a working repository site that includes all the functionality described above. Phase 1 will conclude with a working system for finding, viewing, and submitting subtitles in the hands of hundreds of thousands of Miro users.

Phase 2: Subtitle fetcher in Firefox and on the repository site. Beta of repository site.

In the second phase we will extend the viewing experience to the Firefox browser (for Youtube and non-flash video) and build the alpha repository site into a full-featured site for collaboration. We will develop a Firefox addon that will bring our interface for finding, viewing, and submitting subtitles to every video on Youtube, and to a wide range of non-Flash web video (including any video hosted using the <video> tag). At this point, any subtitled video will also be viewable on the repository site, though only Firefox users will be able to view subtitles while browsing other sites. We will add a set of features to the repository site to make it a full-featured tool for collaboration: users will be able to enter their abilities (familiarity with transcription tools, languages spoken) and interests (genres, etc) in their profiles, and request help from other users in creating subtitles and completing translations. The site will help users to divide tasks and to improve existing translations in response to feedback.

Phase 3: Support more sites / browsers / embedding methods / types of annotation. Build subtitle creation interface.

In the third phase we will extend the user experience to every popular way people watch video, and build-in a simple interface for subtitle creation. We will adapt our Firefox addon to more websites and embedding methods (it needs to be customized to the specific way each site stores or plays video). We will implement the subtitle fetcher

in Internet Explorer and other video players. We can incorporate a simple system for adding subtitles into the viewing experience, both within Miro and in the browser. The specific goals for this phase are best defined as Phase 2 concludes, based on knowledge we gain in the process, feedback from users, and the nature of competitors and collaborators in the space.

4. Other uses and possibilities

This system for subtitling and translation would work just as well for storing and sharing any kind of video metadata, so it has many other potential uses. For example, the repository site could store comments, information about the creator, even information on how to donate to support the creator. All of this could be maintained both by the creator or by users using the same systems that we use for maintaining correspondences between subtitles and videos. As videos propagate to other sites and services, users can help link those new occurrences back to a single page for each piece of content. The repository site beta would be a great starting point for any kind of open voluntary donation system for internet works.

There is already work happening in Miro's open-source community to support time-specific comments and annotations. Two volunteers (from Sao Paolo, Brazil) are implementing support for Youtube-style annotations, as well as an open system (separate from Youtube) for storing annotations. They already have Youtube annotation support working for Miro Linux in testing versions. As we proceed with creating the subtitling infrastructure, we can pursue these other possibilities depending on interest from users, outside organizations, or volunteer developers.

5. Similar projects in this space

dotSUB – dotSUB (dotsub.com) was one of the first web 2.0 subtitling services. Unfortunately, in order to use the service you need to upload your video to the service. In simple practical terms, videos hosted on major sharing sites like Youtube, this presents a challenge for most users; it will be much less likely for subtitles to sprout up organically out of the needs and skills of the audience. In broad structural terms, the dependence on a closed service is political/socially problematic and unsustainable.

Opensubtitles.org – OpenSubtitles.org is an open catalog of user-submitted .srt subtitle files. The service is primarily oriented around mainstream film and cinema (the copyright implications of this are unclear) and depends on a complicated combination of tools. Some architectural decisions (for example, indexing by IMDB#) and its user interface make it non-ideal for the vast majority for user-created video.

Subdownloader – Subdownloader (www.subdownloader.net) is an open-source project that grew out of OpenSubtitles.org. It scans the user's computer for videos, and searches a database (currently Opensubtitles.org) for videos with matching hashes. When it finds a match, it automatically downloads subtitle files according to the user's preferred language. We have been in contact with the SubDownloader team, and there may be opportunities to work together on hash-based searching and on building the repository.

Youtube – Youtube now lets video publishers include captions/subtitles with their programming. However, there is no system for viewers submitting subtitles (except by downloading and re-uploading the same video) or for collaborative translation. This step by Youtube will be important for accessibility and useful for organizations that do subtitling and translation in-house. However, it does not create an open system for subtitling and translating videos (not even on Youtube, not to mention the internet at large). If many publishers begin using captions in Youtube, we can incorporate these captions into our system. And even if Youtube expands their service someday to allow for audience-submitted subtitles, an open system for all non-Youtube video will still be crucial.

Overstream.net / Subtitles.in / Babelbear.com – There is a growing set of services that let people create and view user-submitted subtitles on top of Youtube videos for playback within each site. None of these take a particularly open approach: subtitles are locked within the site. Also, without some kind of browser plugin, there is no way for a user watching a video on Youtube to know that it has been subtitled on one of these sites. We will work with each site to try to get them to open up their platform. Either way, we can make our browser plugin search and display videos from these sites as well (just by searching each site, and loading their own subtitle players).