Liquid FM: Recommending Music through Viscous Democracy⋆

Paolo Boldi, Corrado Monti, Massimo Santini, and Sebastiano Vigna

Dipartimento di Informatica, Università degli Studi di Milano, Italy
`corrado.monti@unimi.it`

Abstract. Most modern recommendation systems use the approach of *collaborative filtering*: users that are believed to behave alike are used to produce recommendations. In this work we describe an application (Liquid FM) taking a completely different approach. Liquid FM is a music recommendation system that makes the user responsible for the recommended items. Suggestions are the result of a voting scheme, employing the idea of *viscous democracy* [3]. Liquid FM can also be thought of as the first testbed for this voting system. In this paper we outline the design and architecture of the application, both from the theoretical and from the implementation viewpoints.

1 Introduction

Most modern recommendation systems use the approach of *collaborative filtering* [8,2]: users that are believed to behave alike are used to produce recommendations. The idea behind Liquid FM is to tip over this approach by making the user responsible for these matches: deciding who they want to resemble becomes a choice of the user, instead of being inferred algorithmically. This scenario can be cast as a voting scheme: each user has to select another one that is believed to be a good recommender. This idea allows us to use this task as a testbed for *viscous democracy* [3].

Viscous democracy is a kind of liquid democracy [6]. In liquid (or *delegative*) democracy, each member can take an active role—by participating directly and exercising their decision power—or a passive role—by *delegating* to other members their share of responsibility. It can be seen as a compromise between representative democracy, where voters are usually neglected any decision making and can only delegate others to do so, and direct democracy, where every voter is called to an active role, regardless of what their inclinations are.

In this sense, liquid voting systems try to take the best from both worlds. Every member's opinion, in a direct democracy, is directly relevant to a final decision, but the vote of each one can be (knowingly!) uninformed; instead, in a classical representative system, elected representatives are encouraged to be informed on the specific decision they are making, but on the other hand the majority of people feel that their opinion on that matter is basically irrelevant.

Liquid democracy permits members to choose among expressing their opinion directly if they feel entitled to do so, or delegating their voting power if they believe others are more capable. Note that these two options are not necessarily exclusive—in our case, in fact, users will be able to do both, if they want to.

Viscous democracy was proposed by Boldi *et al.* in [3] as a particular way to compute the outcome of a liquid democracy voting scheme. It takes advantage of known techniques for measuring centrality in social networks, and in particular it resembles Katz's index [5]. It stems from the assumption that the delegating mechanism should transfer *a fraction* of the user voting power. I.e., if A delegates B and B delegates C, the trust that A puts in C should be less than if A voted C directly. This principle will be further detailed in the next section.

This framework can be used in a variety of settings. In our application, we show how it can be easily adapted to music recommendation. For a certain music genre, we ask users to express a short list of their favorite songs, or to delegate one of their Facebook friends they consider to be an expert on that genre. This builds a graph of delegations for each music genre. We wish to employ this data to create recommendations for each user.

We will detail how we extract information from this graph in Section 2; then, in Section 3, we will describe how we have developed the system: how its algorithms were implemented, the architecture of its components, and the external resources we used; finally, in Section 4 we will sum up our work and present possible directions for future research.

2   Viscous democracy and recommender systems

From now on, we will denote with $d_G(x, y)$ the distance from node x to node y in the graph G, and with $o_G(x)$ the outdegree of node x in G; we may omit reference to G if it is obvious from the context.

Let us define U as the set of users and S as the set of songs[1]. $D = (U, A_D)$, with $A_D \subseteq U \times U$, is the directed graph of delegations; an arc from user u to u' means the former delegates the latter as an expert on the topic. $V = (U, S, A_V)$, with $A_V \subseteq U \times S$, is the bipartite graph of votes, where an arc from $u \in U$ to $s \in S$ means that the user u recommends song s.

We are going to put some restrictions on these graphs: first of all, we are going to assume that there is an underlying, undirected friendship graph $F = (U, E_F)$, with $E_F \subseteq U \times U$, where an edge $(u, u') \in E_F$ expresses a personal acquaintance of u and u'. We impose that $A_D \subseteq E_F$: this permits us to ensure that the trust expressed through a delegation is a result of personal knowledge, as suggested in [3].

Further, we are going to impose that $\forall u \in U$, we have $0 \leq o_D(u) \leq 1$, meaning that a user can delegate only one person, and $0 \leq o_V(u) \leq 3$, meaning that every user can vote up to 3 songs. We are going to consider only nodes

---

[1] As we will explain in Section 3, we are going to consider different sets of songs and votes, one for each music genre treated. For the rest of this section, we are going to consider the music genre as fixed.
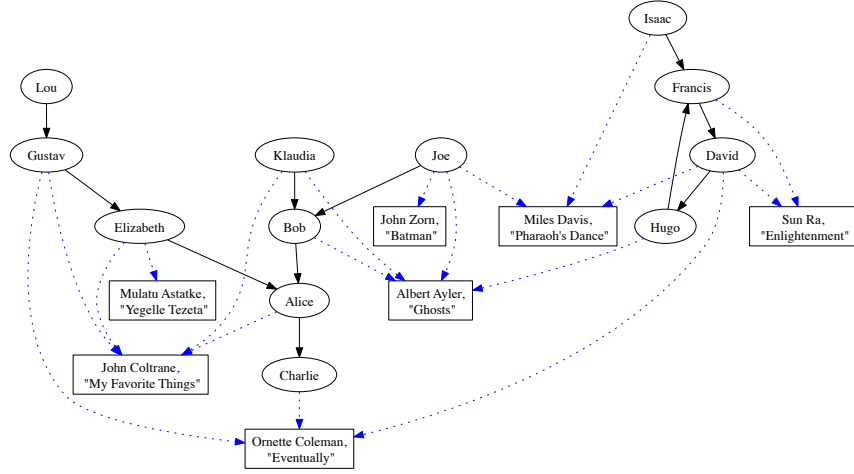
Fig. 1. An example of delegation and voting graphs. Users $u \in U$ are represented with a circle; songs $s \in S$ with a box; the delegation graph D is drawn with black solid arrows, while the bipartite voting graph V with blue dotted arrows.

$u \in U$ having $o_D(u) > 0 \vee o_V(u) > 0$. An example of such a setting is pictured in Figure 1.

2.1 Liquid voting

A voting system is a function $v_D : U \to \mathbb{R}$ assigning a score to each user, depending on the delegation graph. Such a function will be the basic building block of our recommendations.

Usually, in liquid vote this function is just the size of the tree with root in $u \in U$:

$$l_D(u) = \big| \{u' \in U | d_D(u', u) < \infty\} \big|$$

This function is used, e.g., by the well-known `LiquidFeedback`[2] platform. Nonetheless, it assumes that "trust" transferred from a to b is the same whether a delegated b directly, or whether they are connected by a long chain of delegations— and they may not even know each other.

Let us assume that we wish, instead, that the amount of trust passed on from a to b is greater if $(a, b) \in A_D$, and lesser if there are many steps connecting them. To do so, we introduce a *damping factor* $\alpha \in (0, 1]$, defining how much of the voting power of a is transferred to b when a delegates b. Therefore, the

---

[2] `http://liquidfeedback.org/`

$$\alpha = 0.25 \qquad\qquad \alpha = 0.75$$
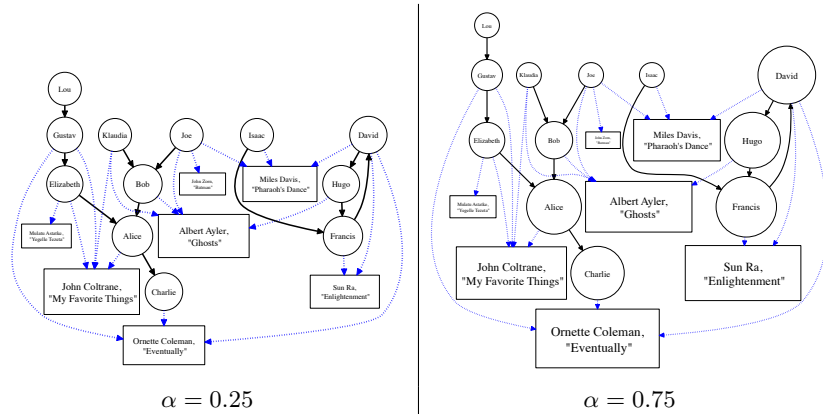
Fig. 2. The same graphs pictured in Figure 1 are here displayed with node size proportional to their viscous score, with two different values for α. Note how a higher α gives higher importance to users delegated by important users. Lowering its value get us closer to a simple vote count. For example, Ornette Coleman's song is ranked higher than Coltrane's only for higher α: this is because it is voted by fewer users, but those are recognized by the community as experts.

scoring function characterizing *viscous* democracy will be:

$$v_D(u) = \sum_{u' \in U} \alpha^{d(u',u)} \tag{1}$$

Authors [3] have noted how, depending on the value of α, the behavior of the voting function greatly differs. For higher values of α, the fraction of trust "lost" in each delegation step becomes smaller and smaller; in fact, for $\alpha \to 1$, we have that $v_D \to l_D$: all the nodes in the tree of u contribute with all their voting power to u, exactly as in pure liquid democracy. Note that if we allow $\alpha = 1$, we must explicitly avoid cycles in D—exactly as with pure liquid democracy; this constraint is not needed with viscous democracy with $\alpha \in (0, 1)$.

With α approaching 0, instead, the voting power becomes nontransferable: all users become equal, regardless of the delegations they received; in other words, the model becomes a direct democracy, without any proxy vote. These differences are presented graphically in Figure 2, making use of the song-scoring function we will show in the next section.

2.2 Global recommendations

Having a score for each user, we can easily score each song $s \in S$. Indeed, we can define a function $r : S \to \mathbb{R}$ as

$$r(s) = \sum_{u \in U \mid (u,s) \in A_V} v_D(u) \tag{2}$$

This function will get us a score for a song proportional to the importance of who voted it, according to v_D. The score is completely defined by the graphs V and D. We can then proceed to rank each song with r, and present them to the users accordingly. As in many standard information retrieval tasks, a user looking for results (about a certain music genre, as we will explain in Section 3) will be presented with all possible items—all songs in S—ranked from higher to lower r. Users will be therefore more likely to listen to songs ranked higher in this list.

Let us call the *influence* of $u \in U$ the difference the votes of user u make in the final rankings—that is, $\sum_{s \in S} r(s) - r_{V \setminus \{u\}}(s)$. Please note that, since we have not normalized r, users giving more votes have a larger influence in the final rankings, serving the purpose of encouraging them to give more recommendations. However, it also explains why we had to put a limit on $o_V(u)$: if we had not, a single user u could have an arbitrary influence on the score r, resulting in the possibility of spam.

In the end, the influence of a user on song scores is determined by the number of recommendations they give—limited, but under their control—and by the delegations they received—unlimited, but not under their direct control.

As mentioned before, an example of how r behaves is pictured in Figure 2.

2.3 Personalized recommendations

The song-scoring function we presented gives the same ranks to whoever is their observer. This behavior is unusual in recommender systems, where the goal is to give the right recommendation to the right person. In our case, a user may be more interested in listening to what their delegate suggested, rather than other—possibly more popular—items. Looking at our example in Figure 2, Francis may be more interested in listening to *"Pharoah's Dance"*, even if it is not globally highly-ranked, because it is the recommendation of his delegate David. Similarly, Hugo may be interested in it, because he has, in turn, delegated Francis.

This goal can be easily expressed as a personalized song-scoring function. Let us define a function $p : S, U \to \mathbb{R}$ as

$$p(s, u) = \sum_{u' \in U | (u', s) \in A_V} \alpha^{d(u, u')} \tag{3}$$

Such a function permits the user u to get a positive score only for the songs recommended by users belonging to the chain of delegations starting in u. For the purpose of maintaining this intention, but at the same time avoiding to completely discard all the songs highly ranked by the original r, we can define a linear combination of the two functions, normalized to 1:

$$c(s, u) = \delta \frac{p(s, u)}{\max\limits_{s' \in S} p(s', u)} + (1 - \delta) \frac{r(s)}{\max\limits_{s' \in S} r(s')} \tag{4}$$

where $\delta \in [0, 1]$ regulates the amount of personalization of c.

An example is pictured in Figure 3.
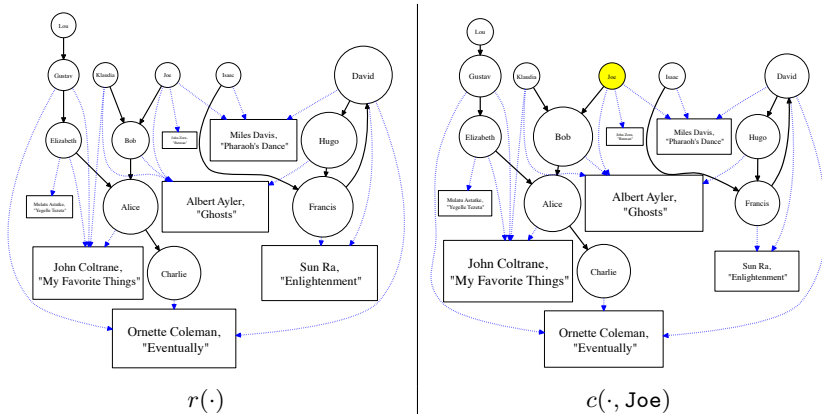
Fig. 3. The same graphs pictured in Figure 2 are here displayed with global song-scoring function r on the left and, on the right, with the personalized function c from the view point of user Joe (in yellow) and $\delta = 0.9$. In the latter, recommendations suggested by the delegate of Joe acquire more importance; those suggested by indirect delegates (namely, Alice and Charlie) increase as well, but by a minor amount.

2.4 Insights for users

In addition to the presented ways to compute recommendations, the setting here described also permits to compute other information that may be of interest to the users. Particularly, it allows them to know how authoritative (i.e., trustable) their taste is in a particular music genre. The function v_D, in fact, can be normalized into a percentile-based scoring, obtaining an easy-to-read assessment in the form "u is better than $\widehat{v_D}(u)$ people out of 100" (for a specific genre), with $\widehat{v_D}(u) = 100 \frac{|\{u' \in U | v_D(u') < v_D(u)\}|}{|U|}$. It can then be used to provide useful information from two different perspectives:

1. Showing to the user a fair evaluation about which music genres they are believed to be more expert about.
2. Presenting to a user interested in learning more about a specific genre which of their friends is considered an expert—making use of the direct knowledge graph defined on page 2.

3   Development

We will now discuss how the presented techniques have been implemented in practice. The final result is Liquid FM: a Facebook application that enable its users to vote one of their friends as an expert on a music genre, and (by means of the described formulas) recommends them some piece of music to listen to, by identifying the best experts.
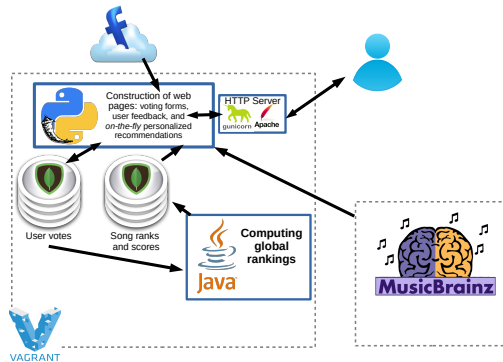
Fig. 4. A schematic representation of the architecture of Liquid FM. An arrow going from A to B indicates data flowing from A to B.

Firstly, we will present a general overview of the architecture of Liquid FM, explaining the role of its main components; then, we will give a more detailed look at the implementations of the formulas presented in Section 2; finally, we will discuss the external components we employed.

Categories As anticipated in the previous section, we applied our scoring algorithms to 9 music genres, called *categories* from now on, and their set will be denoted as C. In this way, we will have different votes and different recommendations for each category. Such a behavior is closer to reality: an expert in HipHop is not assumed to be qualified to give, say, classical music suggestions. However, it also permits to have different graphs for the same users—an interesting fact for future analysis.

The selected categories were Classical, Electronic, Folk, HipHop, Indie, Jazz, Metal, Pop, Rock. They were chosen by inspecting LastFM top 20 tags[3] and discarding those not expressing a musical genre (such as "seen live") and sub-genres (having included Indie and Rock, we discarded "Indie Rock"). We decided to add Classical (only ranked 36th on Last Fm), since it is a different and interesting community, under-represented in services such as the one we referred to.

3.1 Overview

As pictured in Figure 4, Liquid FM features two main components:

– a Java part, with the role of analyzing the whole graphs and computing global scores through v_D and r (equations 1 and 2): it is meant to be fast, and executed periodically;

---

[3] http://www.last.fm/charts/toptags

- a Python part, with the role of glueing the different parts together and providing all the other functions: from the construction of web pages to the implementation of personal scores (functions p and c, equations 3 and 4).

These two parts interact with each other through a shared database, that persistently stores every information. We chose MongoDB, an open-source document-oriented NoSQL database, for various reasons:

- We want fast access in reading and writing data (especially very small chunks, as in delegations and votes) in order to be able to support a large amount of users, like in modern recommendation systems. Moreover, we would like our system to be scalable.
- We want flexibility: since this application is also a proof-of-concept, we need to be able to modify data schemas, totally or partially, without much concern.
- Finally, we do not often need complex operations, involving more than one collection. When it occurs, we would like to control what is happening at application-level, permitting fine-grained handling.

On this database, we have two main collections gathering user-submitted data: following the notation introduced in Section 2, the first stores the graph D and the second the graph V. These collections are both represented in Figure 4 as "User votes". A document in the collection for D looks like this[4]:

```
{    category : c, from : u, to : u'        }
```

While a document in the collection for V has this structure:

```
{    category : c, user : u, advice : s     }
```

The advice s is a dictionary containing author and title of the song, as well as a YouTube video id. In fact, we associate with each song selected by a user a YouTube video, in order to be able to play it as a recommendation. YouTube is in fact one of the largest and most used music streaming platforms, and it can be included in third-party services (with small limitations). A screenshot of the voting phase in displayed in Figure 5.

Please note that the structure of an advice, as well as the category c, is well-incapsulated: therefore, the schemas of these collections can be easily extended in the future in order to support different (i.e., not music-related) scopes.
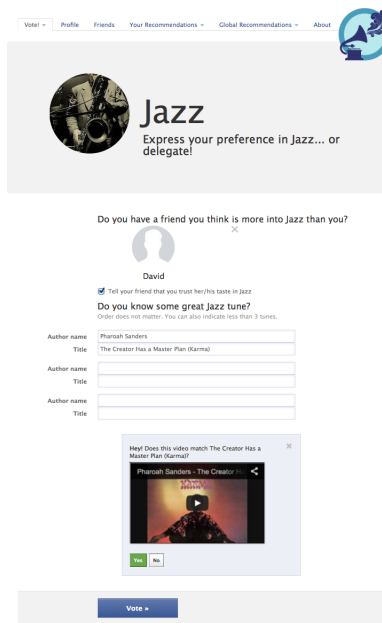
3.2 Recommendations

The division of global and personalized recommendations into two separate components originates from efficiency reasons. Having to compute and store all the personalized scores for each user would be impracticable, as they would

---

[4] Whenever the id of a document is not explicitly expressed, it is automatically generated by MongoDB. This is done efficiently; furthermore, such an id stores the timestamp of creation of the document.

Fig. 5. A screenshot of the voting phase of a user.

be $|C| \cdot |S| \cdot |U|$ scores. Therefore, they are computed with a lazy approach: when a user u asks for her personal recommendations, we compute all of them *on-the-fly* and cache them. Global recommendations, instead, are the main result of the system, and every user depends on them—even to see the personalized scores, since we use the function c (equation 4). For this reason, we compute them periodically with a fast Java component, and save them to a dedicated MongoDB collection.

Global recommendations in Java The global recommendation component was carried out in Java, since for this task it is faster than Python and since efficient open-source libraries to deal with graphs are available; in particular, we employed extensively the WebGraph framework [4] and the `fastutil` library.

This component is run periodically. It takes as input the graphs D and V, memorized in their MongoDB collections, and it results in a new collection for each category, composed of documents of this form:

```
{    advice : s, rank : r(s) }
```

and in another collection ranking users, where each document has this form:

```
{    _id : u,
     category_1: { score : v_D(u), perc : v̂_D(u) }, ...
}
```

We can schematize the process, for each category c, in these steps:

1. Read the graph D from MongoDB and convert it to WebGraph format.
2. Use a parallel implementation of the Gauss-Seidel method (from WebGraph) to compute v_D for each user u. We decided to choose 0.75 as the value of α.
3. Compute the percentile-based normalization $\widehat{v_D}$, and save the user-ranking collection to MongoDB.
4. Read the graph V from MongoDB, identifying the set S of songs. In this step, we also find which YouTube video is the most frequently associated with a certain song, using author and title as identifiers. While doing this, we compute $r(s)$ for each song s.
5. Save $r(s)$ in their collection, indexing documents by decreasing scores. Also save $\max_{s \in S} r(s)$, for normalization purposes.

Personalized on-the-fly recommendations As discussed above, personalized recommendations are computed on-the-fly by a Python component. Python was in fact chosen as the main language of the application, due to its versatility and its fast production times; also, we decided to use `Flask`[5], an open-source web development micro-framework particularly suited for our task.

Personalized recommendations are computed only when users ask for them, since they require to see only a very small part of the graphs, and because storing all of them would be unfeasible. The score we will use to rank personalized recommendations is the function $c(s, u)$ (eq. 4); in order to compute it we must, in the first place, compute p (eq. 3).

To compute $p(s, u)$ for all songs $s \in S$ and a fixed user u we walk through the chain of delegations on graph D, starting from u. Since $\forall u\ o_D(u) \leq 1$, this path on D is unique (although it may end in a cycle). Therefore, we simply proceed as follows (for a suitable stopping threshold ϵ):

---

1. Let `p` be a map with 0 as default value for missing keys.
2. **while** $t > \epsilon$ **and** $\exists u'$ s.t. $(u, u') \in A_D$
     (a) $u \leftarrow u'$ s.t. $(u, u') \in A_D$
     (b) For each s s.t. $(u, s) \in A_V$:
         $\mathtt{p}[s] \leftarrow \mathtt{p}[s] + t$
     (c) $t \leftarrow t \cdot \alpha$

Now we have all the ingredients for function c, and we can proceed to compute the ranking order according to it.

First of all, consider that the ranking order induced by $c(s)$ is equivalent to

$$\bar{c}(s) = k \cdot p(s, u) + r(s) \qquad \text{where} \qquad k = \frac{\delta \cdot \max\limits_{s' \in S} r(s')}{(1 - \delta) \cdot \max\limits_{s' \in S} p(s', u)}$$

[5] http://flask.pocoo.org/

Therefore, for each element s of the map $\mathtt{p}$, we multiply its value by k and add the value of $r(s)$. Then, we retrieve all the other items s s.t. $r(s) \geq \min_{s'} \mathtt{p}[s']$, and insert them in the map $\mathtt{p}$. Finally, we can build the iterator of personal recommendations by chaining two iterators:

1. the iterator of all elements in $\mathtt{p}$, sorted by their values;
2. the iterator of all other elements $s \in S$ having $r(s) < \min_{s'} \mathtt{p}[s']$, sorted by their values; remember that they are already indexed in this order in the database.

The final iterator can be implemented in a lazy fashion, allowing us to retrieve the elements of the second iterator only when necessary. The first iterator, instead, will be computed eagerly upon its request, and then cached. To cache these (and other) values, we employed $\mathtt{redis}^6$, an open-source in-memory key-value cache.

3.3 External services

To conclude this section, we would like to briefly describe the main external software components we used in developing Liquid FM.

Facebook As mentioned earlier, Liquid FM is a Facebook application. The reason for it is that we used the Facebook friendship graph as the graph F defined on page 2. In fact, Facebook is at the moment the largest existing social network (with 1.4 billion users), and it has been previously used as a good approximation of an acquaintance graph [1]. Therefore, we require users to have a Facebook account, in order to limit their choice of delegate to their acquaintances. Accordingly, in the collections described earlier, we used a Facebook-provided id[7] to identify a user u.

Musicbrainz To ensure the validity of the set S of songs chosen by users, we check them against the Musicbrainz database. Musicbrainz is an open music database that anyone can edit. At the time of writing, it contained information about more than 900 000 artists and 14 000 000 recorded songs. Since it follows the open-content paradigm, a user who does not find its favorite song in the database is in principle free to add it; however, the community-review process acts as a filter. Furthermore, Musicbrainz provides a disk image to set up a virtual machine with a fully-functioning Musicbrainz server; we used this approach to be able to access the database fast, without network delays and minimizing the impact on their hosts. Moreover, the database of this virtual machine has been set up to self-update itself periodically, in order to adopt every new edit accepted by Musicbrainz.

---

[6] http://redis.io/

[7] To protect users' privacy, this id is valid only within our app, and cannot be used outside of it.

4 Discussion, conclusions and future work

In this work, we presented a Facebook application aimed at putting the viscous democracy framework [3] to the test. This is at the same time a proof-of-concept of how that voting system can be practically implemented in a real-world social network, and a way to collect data corroborating (or disproving) the supposed advantages of viscous democracy when compared to other, more standard, ways of performing elections in a social setting. An interesting point, here, is that the usage of viscous democracy for recommendation seems to avoid the filter bubble [7], at least in its more algorithmic sense, because this kind of recommendation does not rely on collaborative filtering but is based on a conscious choice. Whether this choice (delegation) can itself induce a similar kind of bubble will be subject of future analysis.

The discussed application is currently active on `http://bit.ly/liquidfm`, and we have so far collected some small datasets; currently, the delegation graphs consist of few tens of delegations, so it is impossible to draw any conclusion from them. In order to be able to collect larger amount of information it is crucial that we find a way to make the application *viral*: this is a matter of social engineering that needs to be taken into careful consideration.

References

1. Backstrom, L., Boldi, P., Rosa, M., Ugander, J., Vigna, S.: Four degrees of separation. In: Proceedings of the 4th Annual ACM Web Science Conference. pp. 33–42. ACM (2012)
2. Bernhardsson, E.: Collaborative filtering at spotify. New York Machine Learning meet-up (jan 2013), `http://www.slideshare.net/erikbern/collaborative-filtering-at-spotify-16182818`
3. Boldi, P., Bonchi, F., Castillo, C., Vigna, S.: Viscous democracy for social networks. Communications of the ACM 54(6), 129–137 (2011)
4. Boldi, P., Vigna, S.: The webgraph framework i: compression techniques. In: Proceedings of the 13th international conference on World Wide Web. pp. 595–602. ACM (2004)
5. Katz, L.: A new status index derived from sociometric analysis. Psychometrika 18(1), 39–43 (1953)
6. O'Donell, G.A.: Delegative democracy. Journal of democracy 5(1), 55–69 (1994)
7. Pariser, E.: The filter bubble: What the Internet is hiding from you. Penguin UK (2011)
8. Ricci, F., Rokach, L., Shapira, B.: Introduction to recommender systems handbook. In: Recommender systems handbook, pp. 1–35. Springer (2011)