

Algorithms Aside: Recommendation as the Lens of Life

Tamas Motajcsek
Gravity R&D Zrt
tamas.motajcsek@gravityrd.com
Hungary

Jean-Yves Le Moine
JCP-Connect
jyleme@gmail.com
France

Martha Larson
Delft University of Technology
m.a.larson@tudelft.nl
Netherlands

Daniel Kohlsdorf
XING AG
daniel.kohlsdorf@xing.com
Germany

Andreas Lommatzsch
Technische Universität Berlin
andreas@dai-lab.de
Germany

Domonkos Tikk
Gravity R&D Zrt
domonkos.tikk@gravityrd.com
Hungary

ABSTRACT

In this position paper, we take the experimental approach of putting algorithms aside, and reflect on what recommenders would be for people if they were not tied to technology. By looking at some of the shortcomings that current recommenders have fallen into and discussing their limitations from a human point of view, we ask the question: if freed from all limitations, what should, and what could, RecSys be? We then turn to the idea that life itself is the best recommender system, and that people themselves are the query. By looking at how life brings people in contact with options that suit their needs or match their preferences, we hope to shed further light on what current RecSys could be doing better. Finally, we look at the forms that RecSys could take in the future. By formulating our vision beyond the reach of usual considerations and current limitations, including business models, algorithms, data sets, and evaluation methodologies, we attempt to arrive at fresh conclusions that may inspire the next steps taken by the community of researchers working on RecSys.

Keywords

Personalization; recommendation engine; machine learning

1. INTRODUCTION

This paper represents the collective thoughts of a group of researchers thinking and discussing what the future of RecSys should be, if allowed to set aside the strengths and weakness of today's algorithms. The group consists of researchers with a spectrum of expertise from RecSys and information retrieval, from senior members of the RecSys community to fresh minds that have just begun their work in the field. Our group cast off the idea that creativity in the area of RecSys must be tied to existing business models, algorithms, data sets or evaluation methodologies. Instead, we pondered how "recommendations in the wild" could support people in their daily and professional lives, both as individuals and as members of their communities. The overarching idea that emerged in the workshop, that we explicitly state here, is that life itself is the original recommender system, teaching people about available options, and providing information that allows them to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

RecSys '16, September 15 - 19, 2016, Boston, MA, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-4035-9/16/09...\$15.00

DOI: <http://dx.doi.org/10.1145/2959100.2959164>

make decisions. The recommender system community can increase the scope of its vision if it temporarily sets aside consideration of the capabilities of current algorithms, and considers how RecSys fall short of serving to strengthen and enhance the lens of life. We believe that the focus must shift back to real life situations, similarly to when social filtering motivated collaborative techniques, to find new applications and approaches for RecSys.

Twenty years ago, the basic idea of RecSys originated from the observation of social phenomena (social filtering). The initial insights of collaborative filtering yielded countless important results and widely used applications. However, we find that RecSys achieving adequate performance do not aspire to high enough goals. We believe that going back to the roots of recommender system research can inspire the research community to develop new techniques that address a number of important and unsolved challenges, and to focus on some new real needs of the users. This paper represents a return to the roots: an effort to again observe life, people's daily activities and social interactions, and to find new potential functionalities, internal system behaviors, and applications for RecSys. We hope that the results of our efforts are useful for fostering the future development of the recommender system field as a whole.

This paper is structured in 3 parts. Firstly, we assess the current limitations and pitfalls of RecSys. Secondly, we ponder the notion of whether "life" could be a model for RecSys, and lastly, we present our ideas about the kind of recommenders that could exist in the future.

2. TODAY'S RECSYS HAS NO CLOTHES

As the story goes, it took a child to stand up and say that the emperor was wearing no clothes. More than one among us was willing to make a similarly naïve but true remark about the state of RecSys today. One author spent three hours interacting with Amazon to find a simple birthday present for her father. Others have had the experience of scrolling through all the rows of Netflix recommendations and still arguing with their spouse over what to watch. As our discussion progressed, more authors became willing to share stories where large commercial recommenders have failed them.

Less easy to elicit are examples where recommendation fails to penetrate domains in which people really could use a helping hand. Classical examples present a limited view of the scope of a recommender: it is not just about Netflix and Amazon, as RecSys have the potential to accomplish much more than list books to read

or movies to watch. We all remember the vision of artificial intelligence becoming a personal assistant that is able to tell us what to do next at any time and in any situation. We have narrowed the scope of our vision quite a bit from there, and if we blink and look again, we can see how our current recommenders are actually stark naked. In Section 4, we will return to RecSys that transcend the current recommendation engines in terms of items and situations in which they recommend.

2.1 People Are Not Machines

Obviously people are not machines. We wake up in the morning and stretch as we get out of bed. Do we find ourselves accidentally mistaking ourselves for deterministic automata? If not, and our impression is that we do not make this mistake, then we should ask the question of how the idea arose that an algorithm that knows only about the past can give us suggestions for the future. Likely it is because we are indeed creatures of habit: we like to have our tea on time, brush our teeth daily, and, depending on age, the bedtime story or nightcap is an essential finish to any day. We also like to hang out with friends, and if we do some of the same things that they do, we have more to talk about. For these reasons, RecSys that look at our past patterns do, in some sense, work. However, if we are not mechanical at our core, then RecSys are well served to abandon the presumption that if they can learn enough about a user, they will know what to recommend in a given situation.

People must have a choice about the kind of recommendations they receive in order to make a true personalized RecSys, instead of just individually-targeted personalized recommendations. And today's RecSys don't seem to be quite up for that task. Imagine a user, Tom, buys an item. That data will forever influence the recommendations he receives. On the other hand, if he deletes it he loses himself in a way. It will be important for Tom to have the possibility to edit and correct the data collected about him, and/or to inform the system about which data are currently relevant.

RecSys should also be able to assess their confidence in their knowledge of user needs. It is not enough to simply output the items that Tom is most likely interested in. When the system is not confident enough in its conclusions it would do much better to wait for user input, or to solicit it by means of well-designed active learning procedures. The cost of a bad recommendation could be enormous in many domains, and a system that is not able to understand its limitations may simply not be acceptable.

Suppose that the RecSys is able to get enough information about Tom in a certain context to supply him with recommendations. Has the system, then, achieved the goal of satisfying Tom's needs? We believe that the answer is no for two reasons. Firstly, telling Tom to choose what he actually already wants has limited value, as such recommendations are often no more than reminders to do something the user would have done anyway. Secondly, the user's choice will interact with the context to create a new context. As such, the RecSys has not provided the user with information on the new set of preferences that are formed as a result of the interaction.

2.2 The Voice Of Recommendations

RecSys differ according to the amount of authority they have in their domain. Reflection on the difference between the word "Recommend" and the word "Command" is informative. It would be ridiculous for a RecSys to start giving "Commands", as people want control over their own actions, and over their lives in general. Navigation system designers bend over backwards trying to figure

out how to tell drivers where to go, while not offending them with a "know-it-all" system.

Today, we face widespread failure of RecSys to support users in consciously reflecting on power relationships. We do not stop to think: does this recommender system "know" more than I do? Can I trust it in everything? Or, should I regard it as a fun Crazy Eight ball: shake it and see where it wants to send me next in order to inspire my own creativity. Instead, it is essential that a user is able to develop different trust relationships 1) with different RecSys or 2) with the same system in different situations where different expertise is required. The level of deference to the authority of the recommender system must be different for a system that proposes medical treatments, as opposed to one that provides lists of movies. Users need to be aware of the power balance between the expertise represented by the recommender and their own expertise.

In order to get there, we need to understand how we can construct a humanlike relationship between RecSys and their users. We see a recommendation as part of a dialogue between persons and technology. As any dialogue, it is an iterated bidirectional process, built upon explicit or implicit, verbal or non-verbal messages ("signs") exchanged between the entities involved (the human and the machine). It involves a progressive interpretation of such messages in order to build mutual understanding, and to foster the expected behavior on both sides. From this perspective, the principles that make a dialogue reliable and effective (as stated in linguistic and semiotic theories) can guide the design of the recommendation process, from the elicitation of users' needs, interests, and intentions, to the construction of the "right" recommendations and their communication to the user.

Trust and verification will be key factors in understanding how users perceive their recommendations. Just because they trust the context, doesn't mean they should automatically go along with the recommendation.

There are some users that want to have their hand held, and others do not want such omniscient advice. We need to proceed on the premise that users' preferences for items only represent one dimension for personalization. Another dimension is the amount of active input that they would like to have from the recommender system. RecSys today lack a user typology regarding expectations about system behavior. Some users may not desire the system to fully solve their problems for them, and perhaps want to follow a more exploratory approach, or gently helped to find their way more autonomously. As such, people will have very different styles in the relationships that they will want to develop with RecSys, and these styles may vary across domains.

There are also ethical limitations. Assuming we have a recommender system that constantly provides us recommendations, the question arises whether that resembles a "Big Brother" who guides us. What personal or private data will be available to the system(s) in order for recommendations to be generated? Will the result be a problem for society if we train ourselves to constantly listen to recommendations?

2.3 When Is Enough Enough?

The assumption that recommendations need to be ranked seems to be a holdover from the printed page. Our reflex is to write things in a list. Just as we start at the top of our to-do lists scribbled with pens and pencils on pieces of legacy paper, we start at the top and expect

to do this or that, and finally arrive at the finish line. Useful RecSys should free themselves from such linearity, and from the idea that there is a best/optimal total order of items that should be determined by the unknown utility function of the user. While some research has been devoted to multi-criteria RecSys, they are still only using multiple criteria to better predict an overall global criterion, which is used to finally rank the recommended items. In reality, people usually consider several factors before coming up with a choice. RecSys should support multiple criteria decision making processes rather than trying to immediately push an integrated ranking based on some machine calculated aggregation strategy.

Another important limitation of current RecSys is that they don't know when to stop recommending. A fine balance is needed to ensure that RecSys are developed to the point where they can understand the needs of the user without making them feel harassed.

Cognitive human limitations must also be taken into account, e.g., what the users want from recommendations, as how many recommendations they want may depend on their mood at a certain moment, and the context they find themselves in. For some, interaction with a RecSys might be a pastime or an end in and of itself. However, in general, the RecSys should not exhaust us: we need plenty of time and cognitive resources left over to enjoy our recommendations, or generally get on with things.

3. LENS OF LIFE

The key driving idea that arose in our process of collective reflection on RecSys is: Life itself is the original recommender system, and people are the query. Without digital devices, algorithms, or data, considering how life itself throws new experiences in our paths and helps us to make choices provides a rich source of inspiration for paths that the recommender system community should take. One could foresee RecSys that mimic life, with its abundance and complexity of messages that can then be interpreted as recommendations to the human senses. In this scenario, the user is the query and life, with its infinity of propositions, is providing a variety of recommendations.

When we stand at a crossroads in life we often seek out those who are wiser or more experienced in order to decide what to do next. This insight is already well represented in existing recommendation systems that combine personalization with the opinion of experts, which is mined from the opinion of the crowd.

3.1 The World Within Us

When considering life itself as a recommender system, two diametrically opposed perspectives arise. On one side stands the perspective that life is a mystery, whose pieces fit together in magical and wonderful ways. This perspective holds that RecSys have the potential to connect people and their expertise in a global "brain", similarly to what happens with neurons in a human brain. As in human brains, we are not aware of how precisely information is produced, stored, and transmitted between neurons through our synapses. Unlike human brains, we are able to have access to this information in RecSys – as such, RecSys may hide from the users the complexity and distributed nature of the wisdom of the crowd.

On the other side stands the perspective that life is intuitive and understandable. People use simple rules to engage with their

surroundings: my big sister is driving and I want to drive too, I became a lawyer because that is what everyone in my family does, and my favorite movie star looks smashing in those shoes so I need to own a pair as well. This perspective holds that recommendations should be perfectly transparent. In fact, a recommendation that is understandable is more valuable to the user than one that is made for unexplained, mysterious reasons.

3.2 The World Around Us

Most of the time the answer presented is obvious, and appears before our eyes. But sometimes life presents the information differently: it produces a synchronicity, something that looks like pure coincidence, but is not. Humans cannot attend to all the information contained in the world around them at once, as we are constrained by our ability to pay attention only to a limited number of things. By objectively assessing the underlying dynamics and intricacies that are hidden from our senses and/or attention, but nevertheless reflected in large amounts of data that can therefore be extracted by analytic means, RecSys could help users find the least probabilistic path to follow in order to attain a better life. RecSys could, perhaps, also present users with answers that are not "superficial or immediate", but are of a deeper or hidden, and of a more "long term" type.

Life doesn't always give us the easiest answer, we know that! Sometimes we have intuitions, but don't know whether they'll serve us well. Sometimes we're not even sure how to follow them! When we become too introspective, we act as bad RecSys, only wanting to select items and achieve goals. How to be aware of the thoughts that come to us and seem necessary to pursue? When you follow your intuition and take this course of action, you find yourself in a place where a synchronicity occurs, and you're given just the right amount of information and clarity you need at that moment.

3.3 The World In Front Of Us

Possibly the most important thing that life tells us about RecSys is that individual experiences (e.g. of achievement or enjoyment), come together to form larger paths of meaningfulness. As such, RecSys should not only support the next act of consumption, but rather infer our future, and help us to build a picture of where we would like to go. To do so, they need to support people by understanding, and helping them understand, how they think. We make decisions in different ways: sometimes we make decisions in an intuition-based snaps of the fingers driven by strong emotions, and sometimes we make decisions slowly and analytically, supported by extensive evidence gathering. To anticipate and support such decision making, the recommender system should take into account the user's emotions, and the different ways in which users make decisions.

Whether we live in a deterministic or non-deterministic world, it often seems as if life knew the future, or at least what is best for us. In a way, these synchronicities show us our best possible future. So, instead of RecSys that only use information about the past, how about a recommender system that could, in a way, infer and thereby "influence" our future?

4. RECOMMENDERS YET TO COME

A central observation that is inspiring our analysis is that today's RecSys are not yet fully user centric. We have mentioned a number of areas to consider: (i) their relationship with the user, including "Command" vs. "Recommend" tone, amount of user involvement, knowing when to prompt for more information, and knowing when to stop recommending, (ii) understanding humans and helping them understand themselves, including working within the limits of human cognition, understanding multiple criteria decision making, and understanding how humans make decisions, and (iii) future-orientedness, including making recommendations based on synchronicity, and inferring and influencing a person's future. In short, RecSys must form a more comprehensive and realistic picture of who their users are and in which directions their needs can unfold in different contexts.

4.1 Moving Us Beyond Items

Moreover, in addition to rethinking the role of recommendations we need to expand the range of items that the recommender will select and present to the user. These systems could provide more diverse types of suggestions to users instead of direct recommendations for items. We believe that eventually, courses of action, intentions, emotions, ideas and even other people could be the "items" of future recommendations. In this sense, RecSys could inspire, assist, decide, and much more, in a completely unobtrusive way.

Following this idea of learning from the richness and unexpectedness of life to build better RecSys, we organized the recommenders we sketched out during our sessions as follows:

Firstly, we came up with the idea of an ultimate RecSys, one that could recommend everything in daily life, a kind of ultimate personal assistant similar to what Spike Jonze showed us in his movie, "Her"; a RecSys which acts as an extension of our consciousness, guiding us, while steering away from becoming a "Big Brother". With access to all our data, this consciousness extension could make recommendations based on synchronicity across every context we experience, suggesting not only items, but ways to regulate emotions, places to search for experts on a topic of interest, how best to manage health issues, or simply what to do next. This type of RecSys may grow beyond a life management tool to develop a personalized relationship with the user, learning not only what items to recommend in what context, but also how to interact with the user over time.

Secondly, we imagined recommenders that could improve the quality of our interpersonal relationships and interactions. Even outside the domain of dating recommendations, such systems could help find the right expert when needed, or kindred spirits to share our thoughts, ideas, and feelings with, filling our social circles with beneficial relationships. Such RecSys would have profound effects on the overall quality of day-to-day interactions, connecting us with individuals with compatible interests, emotional needs, and/or compatible personalities. These recommendations may also be made based on momentary needs: e.g. suggesting the friend within one's network best suited for emotional support during a crisis, or the colleague best suited for brainstorming a specific kind of idea.

Another interesting idea that surfaced during our workshops was that this link could be established between real people, as well as with a fictional character. For instance, the expert that you need in a certain situation could be a character from the TV series you love!

This kind of match could present an intriguing way to enhance our lives by expanding the resources available to us.

4.2 Widening Our Circles

RecSys must be a tool that not only makes it easier for the users to lead their life, but also to lead it in a better way. For example, social interactions could be improved by recommending an expert with whom to collaborate, a friend, or even a possible sponsor. RecSys should also establish meaningful links between people in order to bring groups together, to help improve collaboration. This could represent considerable progress in present times when the collaborative economy is growing stronger and stronger. Meaningful recommendation could therefore increase participation and improve the economy of the multitude.

4.3 Helping Us Help Ourselves

A substantial part of multimedia consumption is done with the goal of mood regulation. Usually, when people think of RecSys and mood, they think of systems that recommend music for a candle light dinner, or videos that will make people laugh. However, mood regulation is the conscious use of media to put oneself in a particular state, which allows you to accomplish something. People use music and videos to relax, but they also use music when they want to stay alert or stay motivated, e.g. on a long car trip or when working out. RecSys of the future will need to focus on tasks and goals: providing resources to assist people in accomplishing what they would like to get done.

How about a RecSys that could engage us to strive to get better, one that could improve our motivation? Recommendations could be a part of the gamification of life. We have to be aware of the purpose of the recommendation, however, because such recommenders have the potential to be abused for manipulative purposes. Recommenders should engage people to improve their lives and achieve their goals, and not to follow a dictator, for instance. This raises questions about the nature of, and the distinctions between, personal and collective engagement: in light of such possibilities we must also think about RecSys from a moral perspective.

4.4 Recommending For Inspiration

Another point, related to the previous one, is that a recommender could help people balance their lives. It could help them as a virtual dietitian, with tips on activities to engage in based on their present mood, or how to use their time between leisure and work better. We thus coined a term: "interspiration", intervening inspiration.

The RecSys could help people define their preferences and in the long run, attain happiness and contentment on their own terms.

5. CONCLUSION

This paper is an output of the "Think-Forward Tank" on the Future of RecSys, which took place during the European Conference for Information Retrieval (ECIR) 2016 in Padua, Italy. The overarching idea that emerged during the workshop is this: life itself is the original recommender system, presenting people with available options, and providing information that will allow them to make the right decisions, and the person themselves is the query. If freed from any current restraints, what could RecSys be, and what could they do?

We believe that RecSys should improve and enhance our daily lives. They must be close to us, in terms of having access to our data and in terms of having a relationship with us. As such, they must also know their boundaries, when to recommend and when to stop, and how much interaction is desired. With the enormous amount of data that everybody produces on a daily basis, from the more personal small data (bank transactions, calendar appointments, emails), to the medium data (Netflix’s user profiles), to the Big Data (the vast amount of information on digital social interactions accumulated on Facebook’s data servers), the need arises to give users a stronger sense of ownership. By filtering data based on relevance to context and/or task, RecSys could use the global brain, the “noosphere”, so users feel smarter, cognitively lighter, can observe synchrony, and therefore make better decisions. We also believe recommenders must be in accord with our inner self, recommending what is truly good for us. They should be aware of the needs and preferences of their users, not only to help them achieve their personal (or professional) goals, but to inspire them to set higher goals they haven’t thought about themselves.

Man is the query. RecSys of the future must provide users with answers that not only make their lives easier, but also make their lives better. The recommender system community can increase the scope of its vision if it, temporarily, sets consideration of the capabilities of current algorithms aside, and considers how RecSys fall short of serving to strengthen and enhance the lens of life.

6. ACKNOWLEDGEMENTS

The work leading to this paper has received funding from the European Union’s Seventh Framework Programme (FP7/2007-2013) under CrowdRec Grant Agreement n° 610594.

The third author is also affiliated with Radboud University Nijmegen.

7. ADDITIONAL AUTHORS

Additional authors: Omar Alonso (Microsoft, email: oralonso@gmail.com), Paolo Cremonesi (Politecnico di Milano, email: Paolo.Cremonesi@polimi.it), Andrew Demetriou (Delft University of Technology, email: A.M.Demetriou@tudelft.nl), Kristaps Dobrajs (JCP-Connect, email: Kris.Dobrajs@jcp-connect.com), Franca Garzotto (Politecnico di Milano, email: Franca.Garzotto@polimi.it), Ayşe Göker (Robert Gordon University, email: A.S.Goker@rgu.ac.uk), Frank Hopfgartner (University of Glasgow, email: Frank.Hopfgartner@glasgow.ac.uk), Davide Malagoli (Moviri SpA, email: Davide.Malagoli@moviri.com), Thuy Ngoc Nguyen (Free University of Bozen-Bolzano, email: Ngoc.Nguyen@unibz.it), Jasminko Novak (University of Applied Sciences Stralsund and European Institute for Participatory Media, email: Jasminko.Novak@fh-stralsund.de), Francesco Ricci (Free University of Bozen-Bolzano, email: fricci@unibz.it), Mario Scriminaci (Moviri SpA, email: Mario.Scriminaci@moviri.com), Marko Tkalčić (Free University of Bozen-Bolzano, email: Marko.Tkalčić@unibz.it), Anna Zacchi (Doxee, email: AnnaxZacchi@gmail.com)