

Community Engagement Triage: Lightweight Prompts for Systematic Reviews

MARTIN DITTUS, Nokia Bell Labs

LUCA MARIA AIELLO, Nokia Bell Labs

DANIELE QUERCIA, Nokia Bell Labs

Among platform designers there is still little actionable knowledge about how to foster community engagement, and new community platforms are likely to fail. Much research has documented effects that could inform better designs, however there is a lack of practical methods to review particular concerns. We design a lightweight method for systematic reviews of community engagement scenarios, based on a catalogue of 59 techniques, presented in the form of method cards, and co-developed with community engagement experts. We validate it in problem-solving workshops with organisers, and find that a full review is feasible in under 15 minutes, that it provides effective prompts and important guidance, and can successfully support process innovation. Finally, we validate it for use in summative evaluations of online communities. In a large-scale observational study of 1 million Flickr users, application of the method reveals the surprising impact of a major site redesign: it successfully increased content engagement, but at the expense of social interactions.

CCS Concepts: • **Human-centered computing** → **Collaborative and social computing design and evaluation methods**; *Empirical studies in collaborative and social computing*; Empirical studies in HCI;

Additional Key Words and Phrases: Community Engagement; Online Communities; Evaluation.

ACM Reference format:

Martin Dittus, Luca Maria Aiello, and Daniele Quercia. 2017. Community Engagement Triage: Lightweight Prompts for Systematic Reviews. *Proc. ACM Hum.-Comput. Interact.* 1, 2, Article 39 (November 2017), 22 pages. <https://doi.org/10.1145/3134674>

1 INTRODUCTION

New offline and online communities are likely to fail: it is not enough to build the platform and expect that a community will come [4, 13, 18, 28, 36]. A wide range of research has shown that to increase community engagement, designers also need to be mindful of the affective structure of their platform. This includes processes of affiliation and interpersonal attraction [25, 26], which in turn can help regulate group behaviour [10]. It also includes the degree to which the contribution experience reflects the personal motivations of participants: intrinsic motivations can provide triggers for action, and yield rewarding experiences that can sustain contributor engagement for long periods [9, 37, 39, 41]. In volunteer communities, participation is further constrained by the capacity and ability of their contributors. Some of these constraints can be addressed with interventions, for example by removing barriers, nurturing novices, and fostering a belief that participation is possible and will be welcomed [6]. A wealth of empirical studies has identified particular community engagement strategies that address these concerns, yet they are not always easily accessible for practitioners.

In the academic community engagement literature, there are some attempts to collect existing empirical evidence of important effects, and make the knowledge more accessible to practitioners.

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proc. ACM Hum.-Comput. Interact.*, <https://doi.org/10.1145/3134674>.

Most recently, a publication on building successful online communities collects insights in psychology, economics, and other social sciences, and translates those insights into so-called ‘design claims’, or generalised considerations for system designers [24]. A system designer who is faced with a pressing or complex engagement scenario may not find the time to consult references, especially if they are not already familiar with the underlying theory. This leaves a methodological gap: could the knowledge become the basis for a lightweight method for systematic reviews?

To address this concern, we make three main contributions. First, we present a lightweight process for applied problem-solving by triage: the systematic assessment and prioritisation of problems as a means of deciding suitable next steps. In a 5-step process, we assemble a *catalogue of 59 distinct community engagement techniques* (CETs) in both online and offline scenarios, starting with a systematic literature review of existing CET collections (21 of which come from the foundational work in [24]), and involving reviews by community engagement experts. Based on the catalogue we develop a set of *method cards*, tangible artefacts intended to support group processes [42].¹ The cards are intended for use in card sorting exercises for an individual or a group who seek to reason about community engagement scenarios. We propose *simple exercises* for the use of the card set in formative evaluation (reflection about specific problem scenarios) and summative evaluation (reflection about overall community health). Second, we validate the card set in *problem-solving workshops* with organisers of offline communities, covering a diverse range of concerns in the context of meet-up groups, neighbourhood community initiatives, political youth organisations, and other settings. We find that a full review is feasible in under 15 minutes, that card sorting exercises provide effective prompts and important guidance during problem-solving, and that they can successfully support process innovation. Third, we further validate the card set in a large-scale *observational study* of 20,000 Flickr groups, involving more than 1 million participants. We find that the method cards can easily be translated into measurable indicators as a basis for quantitative evaluations. Our study reveals the surprising impact of a major Flickr redesign in 2013: it successfully increased content uploads and content engagement, but at the expense of social interactions among users.

We will first define our key terms, and describe our process for designing and validating the catalogue. In a closing section, we discuss study findings with practical implications for the use of the CET catalogue.

1.1 What is a Community?

In the academic literature, there is no single accepted definition of the term *community*. In urban planning and public participation literature, a community is understood as a geographically bound social group, often a neighbourhood [30]. In online settings, communities can be understood as networks that are bound by social ties or shared affinities, for example in the form of knowledge-sharing networks, or communities of practice [17].

For the present work, we understand community broadly as groups of participants coming together for a shared purpose, as voluntary members of a *community organisation*. Important properties of such organisations are the presence of a shared identity among its members, a shared organisational purpose, and a degree of openness: participants do not always have to be formal members to be able to contribute, but they can become members at a later time. Similarly, not all members need to be active contributors [34, 35, 38].

¹Available for download at <http://ce-triage.org>.

1.2 What is Community Engagement?

In the context of our method, we understand *community engagement* to be broadly the act of participating in a community organisation. This includes participation in shared practices, and contributing to shared output. For some community organisations, social interactions can form an important element of organisational practice. In larger organisations, participants may engage in the facilitation of coordination processes such as task allocation, evaluation of outcomes, and conflict resolution. They may also engage in the mentoring of first-time participants, for example through the provision of clear guidance, and by making suggestions for desired contributions.

1.3 What is a Community Engagement Technique?

We define *community engagement techniques*, or CETs, as methods to foster community engagement, for example as catalysts for community action. CETs can be applied by individuals or groups in the context of a community organisation. Some CETs increase participant investment in the organisation, and can serve to recruit, train, and retain community members. Others describe methods to regulate community output and interactions, such as the provision of feedback for certain contributions, or means of dealing with toxic members.

The term CET was chosen in reference to *behaviour change techniques* (BCTs) in medical and behavioural sciences. A large body of published research has established comprehensive discourse of BCT, and to date multiple BCT catalogues have been published in the form of taxonomies or frameworks [1, 20, 32, 33]. Our definition of a CET is derived from a definition of BCT published in the literature [32]. Importantly, CETs describe the *active ingredients* of an engagement intervention: the ‘what has to be done’, rather than the ‘how to do it’. For example, a CET may recommend to publish a list of open tasks, to address people individually when requesting contributions, or to provide mentors for newcomers. However, it does not prescribe the specific steps to deliver these interventions.

In our formal definition, CETs are methods that:

- (1) aim to foster community participation,
- (2) are active ingredients of interventions,
- (3) can be used alone or in combination with other CETs,
- (4) are observable and replicable,
- (5) can have a measurable effect on specified forms of engagement.

2 METHOD FOR CREATING A CET CATALOGUE

Initial reviews of the literature surfaced two early collections of engagement techniques for online communities that were discussed in the introduction [24, 38]. We sought to systematically identify more such collections across a wider range of domains, to aggregate their techniques into a larger catalogue, and to present the catalogue in the form of method cards. A schematic diagram of the overall process is shown in Figure 1.

2.1 Step 1: Systematic Literature Review

The initial review took place in two parts, a systematic literature review, and consultation with international domain experts to identify further sources. Our literature review spanned a wide range of domains, including research of online communities, peer production, open collaboration systems, and volunteering. Research in other domains with related concerns was also considered, including literature on public participation in policy and urban planning, campaigning and marketing, and collective action.

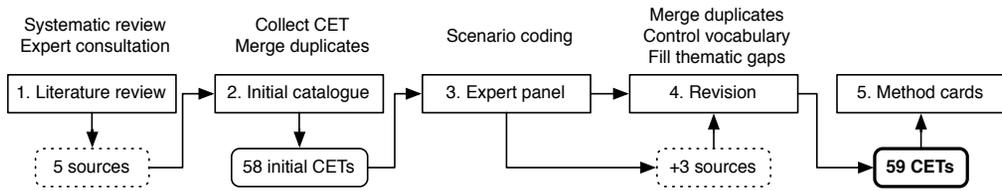


Fig. 1. 5-step process for creating a CET catalogue and card set.

We chose Google Scholar as primary corpus for the literature review. Prior reviews of behavioural and medical studies showed that it has comparable recall and precision to domain-specific corpora [7, 8, 14, 43]. Our search strategy was based on an initial expectation of the theoretical concerns. We sought to cover a wide range of aspects, including considerations of self-selection effects, the importance of social motivators, means of newcomer support, and related concerns. We refined the search strategy in an iterative manner, seeking to increase recall of sources that were already known, and refining our search terms based on reviews of search results. We searched the full corpus for all available years, and excluded patents.

The final search query was structured to include the type of publication, the observed setting, and the observed type of engagement. It is as follows:

```
intitle:(framework OR taxonomy OR classification
OR ladder OR model OR survey) AND
(social OR group OR community OR crowd OR
collective OR connective OR public OR member OR
volunteer) AND
(engagement OR participation OR mobilisation OR
intervention OR interaction OR churn)
```

We then reviewed the search results using a defined set of inclusion criteria. We sought to identify frameworks or survey papers (collections) written in English that provide identifiable lists or tables of specific interventions to foster participation in the context of an organisation, community, or social network, informed by empirical studies. We further limited the review to settings where participation is voluntary.

The following thematic categories of literature were excluded from our systematic review because they did not meet our definition of a community:

- Studies of closed groups, for example where participation is a result of employment. This includes productivity studies of work groups (teams of employees).
- Studies of groups where participation is by circumstances outside the participants' control, rather than by choice. This includes treatment adherence studies in healthcare, and bonding and welfare studies of adolescents in schools.
- Studies focused on individual behaviour change, rather than behaviour in the context of a group or community. This includes workplace health interventions targeted at the individual (rather than employee-organised health groups).
- Studies of settings that exceed the scope of an organisation (unless an express focus is placed on local groups), including political participation in societies, public order and policing in cities, public health studies, economic and social development studies in developing regions, and any domain involving a wider public that is not limited to a specific setting.

- Studies that focus on other concerns, for example learning outcome studies in online learning (rather than studies of social interactions in online learning), or process models in public participation (rather than behavioural interventions).

During the systematic literature review we considered 313 references. 64% of the search results were imminently not suitable based on their title alone, 21% were found unsuitable after manual review of the content, and a further 12% proposed theoretical models or process models rather than intervention strategies. 6 sources discussed engagement effects, but did not provide clear intervention recommendations for systems designers. We identified 3 sources that satisfied the inclusion criteria [27, 38, 40].

To further augment the review, we consulted 8 international experts in online communities and community engagement (4 researchers and 4 practitioners). Recommendations by these experts yielded 10 further sources, 2 of which were suitable for inclusion [16, 24]. In total, 3 papers and 2 books satisfied the inclusion criteria.

2.2 Step 2: Initial Catalogue

For this first version of the catalogue, we sought to identify techniques that can generalise to a wider range of community settings. We reviewed the 5 sources identified during the literature review to identify specific recommendations for interventions that can be expressed as CETs.

For every source, we identified summary tables or lists of engagement techniques that could be used to structure interventions. Among those, we identified techniques which fit our formal definition of a CET. Such techniques must aim to foster community participation, they describe the active ingredient of an intervention (the *what* rather than the *how*), they must be observable and replicable, and they must have an empirical basis. We excluded references to theoretical models, motivational categories, and other concepts that do not refer to an explicit intervention action.

In total, we collected 77 community engagement techniques from these sources. We then reviewed the collection of techniques in an iterative manner to identify duplicates. These are instances where multiple sources referred to the same intervention. For example, two sources provided evidence that requests can be more effective if they are made by a person who is trusted by the recipient, or by a person in high social standing. One source recommended “Requests from people who are attractive, have high status, or the requestor likes, is familiar with, or is similar to” [24], the second suggested “Encouragement by friends, family, respected authorities, advertising” [38]. We merged such instances, taking care to preserve their active ingredients.

This yielded an initial list of 58 distinct CETs. We summarised each with a short sentence, using active verbs to place an emphasis on recommended organiser actions. We further included important constraints and caveats where these were provided by the literature. For example, several sources recommended to credit groups for shared achievements, rather than crediting individuals, and one source provided specific evidence that comparative individual feedback can foster competition among participants. We summarised this as follows: “Credit groups for shared achievements. Only give comparative individual feedback if you seek to foster competition”.

2.3 Step 3: Review by Expert Panel

This catalogue underwent multiple iterations to test for phrasing, thematic coverage, and relative weighting of topics. Initial iterations were executed by the three authors, this was followed with a review by a panel of four community engagement practitioners. All panelists were asked to code one or more community engagement scenarios of their choosing, using the catalogue of techniques to describe an engagement concern or potential intervention. To this purpose, they were asked to scan the catalogue of techniques, and identify techniques that could be used to describe the chosen

scenario. They were then asked to highlight any thematic omissions, and to suggest alternative phrasings where language was ambiguous.

This yielded several recommendations for improvement:

- Several panelists identified more instances where duplicate techniques mapped to the same underlying active ingredient.
- One panelist recommended to use a controlled vocabulary for key terms to improve consistency. Two panelists recommended to simplify the language, for example by removing technical terms, and to express techniques as suggestions rather than prescriptions.
- One panelist observed a thematic emphasis on top-down techniques, for example a relative emphasis on reward systems. To broaden thematic coverage, the panelist recommended to include techniques related to group self-governance and bottom-up processes.

The review further revealed that the catalogue does not currently address certain strategic concerns, including means to establish fundamental rules of interaction, negotiate priorities, or select organisational structures. This is an artefact of its grounding in empirical behavioural sciences, where studies focus on effects in individual and group psychology, rather than questions of organisational culture and purpose.

2.4 Step 4: Revised Catalogue

Several important changes were made in response to reviewer feedback. Duplicate techniques were merged. To improve consistency, we introduced a controlled vocabulary for recurring terms, for example by replacing the terms leader, administrator, and moderator with the term ‘organiser’. We simplified the language, avoiding domain-specific or technical vernacular such as ‘legitimate peripheral participation’.

We identified three further sources to address an observed shortage of techniques relating to emergent and group-led practices. Two were identified with a manual search for literature on open collaboration [3, 12], and one through a review of reverse citations of already identified sources [23]. This helped address a comparable lack of techniques relating to individual autonomy, community outreach, and group cohesion.

The revised catalogue is listed in Table 1. It contains 59 unique techniques, based on 110 techniques identified in the literature. For ease of access and review, we arranged it into six domains of community activity, chosen based on the primary intended outcome of each technique: discovery of tasks, appeals for contributions, techniques to support participation, evaluation of outcomes, social encounter, and collaboration among participants.

Among the 59 techniques, 43 are unique contributions attributable to a single source, the remaining 16 are derived from two or three sources. The largest overall contributor provided 25 techniques, 13 of which are unique contributions [38]. The second-largest contributor provided 21 techniques (16 unique) [24]. Both are focused on online communities. In comparison, 24 techniques are found in sources focused on offline communities (10 unique) [16, 27]. In other words, a single reference source alone could not have supported a catalogue of the same thematic breadth.

Domain	#	Community engagement technique	Sources
Discovery	1.	Promote repeatedly across a range of media platforms to ensure visibility	[16, 38]
	2.	Appoint a person responsible for recruitment	[16]
	3.	Build relationships with prospective members	[16]
	4.	Present community contributions well and update them frequently to encourage repeat visits	[38]
	5.	Make it easy for people to search and discover community contributions that interest them	[38]
	6.	Maintain a list of requested contributions to increase the likelihood that someone will provide them	[24]
	7.	Allow members to track new and ongoing work, for example through opt-in notifications	[24]
	8.	Define the topic of the group well to recruit members with particular interest in the topic	[23]
Appeals	9.	Keep requests simple to avoid early loss of interest	[24]
	10.	Emphasise the benefits of contributing to motivate those who care	[24]
	11.	Appeal to shared fears if necessary, but carefully explain your reasoning	[24]
	12.	Ask others to issue requests on your behalf, especially those who are well-regarded by the intended audience	[16, 24, 38]
	13.	Show that others are already contributing	[24]
	14.	Issue specific and highly challenging goals	[24]
	15.	Set concrete deadlines to increase the likelihood of community contributions	[24]
	16.	Highlight people's individual ability to make unique contributions	[24]
	17.	Tailor requests to people's interests and capabilities	[24]
	18.	Address people individually instead of broadcasting requests	[24, 27]
	19.	Recruit people who have prior relationships outside the group, this can yield stronger ties between members	[23]
Participation	20.	Provide clear guidance and navigation to avoid confusion	[38]
	21.	Issue clear norms and policies about acceptable conduct and appropriate contributions	[23, 38, 38]
	22.	Foster empathy, trust, mutual support, and shared pursuit	[38]
	23.	Provide newcomer support and training, for example in the form of tutorials and demo videos	[16, 27, 38]
	24.	Remove technical, logistical, and physical barriers to entry	[3, 38]
	25.	Provide a safe space for members, and be considerate of their privacy	[27, 38]
	26.	Provide simple starting points so that people can start gradually with small contributions	[38]
	27.	Provide support for large and frequent contributions by highly active members	[38]
	28.	Design immersive experiences with hard challenges, but leave members in control, and provide clear goals and feedback	[24]
	29.	Give members strategic autonomy in their work	[16]
Evaluation	30.	Provide tools and processes to deal with spam and vandalism	[38]
	31.	Implement effective processes to resolve differences and disputes, and deal with unhelpful members	[38]
	32.	Engage people in reflection about their work to foster gradual improvement	[16]
	33.	Ask members to review contributions and provide constructive feedback	[38]
	34.	Highlight contributions that are complementary, and emphasise the collective undertaking	[24]
	35.	Give feedback in relation to specific goals, in particular when it is positive and sincere	[24]
	36.	Issue tangible rewards for good contributions, but be mindful that this can foster competition	[24]
	37.	Issue tangible rewards for simple tasks to increase contribution volume, but carefully observe contribution quality	[24]
	38.	Avoid tangible rewards based on contribution volume if quality is a concern, it encourages gaming the system	[24]
	39.	Make member contribution activity visible to others to foster interaction and collaboration	[3, 16, 38]
	40.	Provide recognition and respect for high-quality contributions, and for providing guidance to others	[38]
	41.	Credit groups for shared achievements. Comparative individual feedback can foster competition	[3, 24, 27]
	42.	Recognise different kinds of contributions, emphasise that there are many ways to participate	[27, 38]
	Encounter	43.	Bring together disparate groups, and encourage different points of view to improve understanding of a complex task
44.		Give everyone an opportunity to speak, prevent individuals from dominating meetings and activities	[27]
45.		Develop a common language to foster mutual understanding and a shared identity	[27]
46.		Implement ideas that emerged in the community to show that shared discussions can have a real impact	[27]
47.		Encourage social contact between members to increase enjoyment, and make tedious tasks more engaging	[16, 24]
48.		Offer a variety of ways for members to get to know each other, as different people like different modes of encounter	[27]
49.		Provide ways for members to identify relevant and competent collaborators	[38]
50.		Avoid anonymity or aliases if you want to discourage social loafing and free riders	[23]
51.		Highlight commonalities among members to foster familiarity, but avoid excluding others	[3]
52.		Foster a sense of community and shared identity among members	[16]
Collaboration	53.	Maintain cohesion by carefully controlling group size: set growth limits, cull inactive members, form subgroups, ...	[16, 23, 24]
	54.	Nominate mentors, and let experienced members guide newcomers	[27, 38]
	55.	Promote charismatic organisers with visionary goals to inspire participation	[38]
	56.	Celebrate members by highlighting their efforts, providing tribute, sharing historical narratives	[38]
	57.	Give organisers special powers, but state your expectations and hold them accountable	[27, 38]
	58.	Include a broad range of members in important decisions to benefit from a diversity of perspectives	[23, 27]
	59.	Provide support for collaboration: communication, idea development, task management, peer support, ...	[3, 38]

Table 1. Catalogue of community engagement techniques.

2.5 Step 5: Translation to Method Cards

We converted the catalogue of techniques to a set of method cards, tangible artefacts intended to support group processes [42]. This choice was informed by prior examples of card sets that translated generalised knowledge catalogues into tangible artefacts, including IDEO method cards [19], the Arup Drivers of Change card set [29], and the Design Heuristics card set [21]. The physical format of these cards makes them highly suitable for in-person workshops. Method cards can be used in

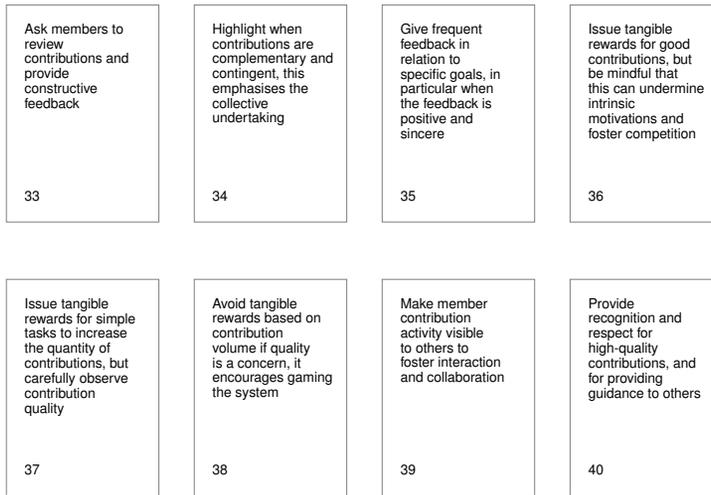


Fig. 2. Card design used for validation studies.

card sorting exercises and similar review processes, and in group settings the cards can become important communication tools [42]. A visual example of the cards is shown in Figure 2. The cards are available for download at <http://ce-triage.org>.

3 EXERCISES FOR COMMUNITY ENGAGEMENT TRIAGE

The catalogue of CETs is not an evaluation model in itself, but it can assist in the evaluation of communal processes. Using the catalogue as a reference, organisers can review techniques to foster certain participation activities, and identify participation activities that are not currently well-supported. Towards this purpose, we propose two broad uses of the catalogue: in *formative evaluation* to assist in problem-solving of specific scenarios, and in *summative evaluation* to provide an overall health check of a community platform or organisation. We present instructions for both uses. These exercises were chosen to illustrate the versatility of the catalogue, and they should be regarded as exemplary uses, rather than fixed instructions. Practitioners are invited to adapt them for their own purposes. In later sections we present two validation studies of the use of the catalogue in evaluation settings.

3.1 Formative Evaluation

A first set of exercises is designed to assist in formative evaluation during the earlier stages of building a community, with a specific focus on problem-solving. Their aim is to systematically reflect about a stated concern, while broadening the range of considerations of contributing factors. They were designed to provide means of discovery and reflection, rather than trying to offer mechanistic solutions to complex scenarios. Their outcome is a set of actionable hypotheses about potential root causes, and suggestions for specific interventions. Practitioners can then derive specific design recommendations based on the underlying principles provided by the CET catalogue. We test formative evaluation in our first study, and provide an applied example at the end of our second study.

3.1.1 Closed Card Sort. In a closed card, the CET method cards are sorted into a set of pre-defined categories, identifying a subset of CET that relate to the problem at hand. This can be a good initial step in problem assessment.

- (1) Describe a community engagement problem that you have experienced.
- (2) Review the deck of cards, one at a time.
 - For every card, consider whether the concept described by the CET might relate to the stated problem.
 - Sort the cards into three stacks: CET that definitely apply to the problem, those that might apply, and those that definitely do not apply.

3.1.2 Open Card Sort. In an open card sort, cards are manually grouped according to emergent themes, without a reliance on previously defined categories. This can be a useful next step after a closed card sort, for example to review CETs that were identified as relevant to the concern, and organise them into a set of broader topics.

- (1) Describe a community engagement problem that you have experienced.
- (2) Decide on a purpose for your review. For example, to identify potential root causes of the stated concern, or to identify opportunities for potential interventions to address the concern.
- (3) Review the cards, one at a time.
 - For every card, consider whether it links to a broader theme.
 - Sort the cards into stacks corresponding to these emergent themes.
 - If your understanding of the scenario changes during this review, reorganise the stacks to reflect your new understanding.

3.1.3 Combined Card Sort. For more experienced reviewers, the two exercises can be combined: while reviewing the cards, place those aside that definitely do not apply to the stated concern. Organise the remaining ones into emergent themes as you sort through the stack.

3.2 Summative Evaluation

A further exercise is designed to assist in summative evaluation. It is not driven by specific concerns, but instead provides a broader overview of the overall health of a community system, with a focus on community engagement aspects. We present its use in quantitative evaluation, and test it in our second study. Practitioners may alternatively prefer qualitative forms, collecting specific experiences and outcomes across a wide range of activities and concerns.

- (1) Declare a scope for the evaluation. For example, to assess which affordances of the platform foster particular forms of community engagement.
- (2) Identify an evidence base for your assessment. For example, data about participant activity as recorded on a web platform.
- (3) Review the method cards, one at a time.
 - For every card, consider whether it applies to the declared scope. Put others aside.
 - Every CET provides a recommendation for an organisational activity. Discuss how you might measure your adherence to it.
 - Based on this, develop one or more indicator measures for every CET.
- (4) Group indicator measures into thematic categories.
- (5) Compute all indicator measures, provided the available evidence allows this.
- (6) Review the findings.

The findings of summative evaluation can provide important input for strategic development planning. They can be evaluated in relation to prior expectations or aspirations, and observed over

time. For example, summative evaluation can identify platform affordances that have important engagement functions. In some cases, a review of the findings may reveal ongoing engagement problems, or opportunities for the development of new activities. Platform designers may then proceed with a formative evaluation stage for every individual concern, acquiring a better understanding of it, and developing them into design recommendations.

4 FIRST VALIDATION STUDY

4.1 Problem-solving by Domain Experts

We subsequently assessed the use of the catalogue to support community planning and evaluation. In a first stage, we tested its use in formative evaluation to support problem-solving by community organisers. This first study had multiple objectives. Most importantly, we sought to determine whether the catalogue of techniques can be made useful for practice, whether it could provide suitable guidance, and whether it could support process innovation. We further sought to assess whether the relative size of the catalogue posed a practical barrier. How long does it take to review the 59 techniques? Is it possible to shorten this time investment when needed? Finally, we sought to identify existing organisational processes where the catalogue could be of use to organisers, and assess to what extent facilitator help was needed to aid in the interpretation of the techniques.

4.2 Card Sorting Exercises

The cards were used in 5 problem-solving sessions with community experts, with a focus on different card sorting exercises. All sessions were timed. Two longer sessions were held with professional community organisers. Sessions lasted between 2-3 hours and were held in the homes of participants. A first session involved an organiser for meet-up groups in online and offline settings. A second session involved two organisers of local initiatives to improve community wellbeing. Three shorter sessions with four community organisers were held in a social setting and lasted between 10 minutes and 1 hour. A first session involved a theatre-maker and arts educator, a second session was shared by an expert in community engagement for science communication and a theatre director and facilitator, and a final session involved a former organiser in a political youth organisation. Across these sessions, only one expert stated that he currently makes use of a systematic process when addressing community engagement problems.

All participants were first given an introduction to the card set, and started the session with a closed card sort. They were asked to describe a community engagement problem they had experienced, and to discuss any factors that may have contributed to it. They were then asked to review all 59 cards, sorting them into one of three groups: 'definitely applies to my scenario', 'maybe applies', 'definitely does not apply'. After the card sort, they were asked to discuss to what extent their understanding of the problem had changed. They were further asked whether the cards yielded plausible suggestions for potential remedies (they provided *guidance*), and whether they yielded previously unconsidered ideas (they supported *innovation*).

During the two longer sessions, this was then followed with additional exercises. Cards that had been deemed applicable to the stated problem were reviewed in a second open card sort, with the intention to identify logical groupings between them. Figure 3 shows the outcome of such a card sorting exercise during one of the problem-solving sessions. The sessions were concluded with an open conversation about the nature and design of the cards, including a discussion of their potential uses and limitations.



Fig. 3. Example of an open card sort, where method cards are arranged in thematic groups.

4.3 Results

In these sessions, focused individuals could work through the full set of 59 cards in 15 minutes, provided they did this in silence. One participant commented on every card, this significantly slowed down his pace. He chose to interrupt the review after 25 minutes, at which point he had only reviewed 31 cards.

According to participants, the techniques yielded *plausible suggestions* that were suitable for the problems at hand, and that could easily be linked to their practice. One expert stated: “For every card I thought: yes, that’s how it’s done. However I’ve inherited a lot of this, and don’t always know the reasons.”

Card sorting exercises yielded recurring moments of *innovation*. In each of the three longer sessions, participants were able to derive previously unconsidered ideas from the exercise, even when they were already familiar with the individual techniques. Sometimes this was achieved by a shifting of perspectives while reviewing cards. As one expert stated, “Some of these really helped reframe the issue”. One expert characterised the process as an effective catalyst for reflection: “Really good! This is making me think.”

When discussing potential uses of the catalogue, several participants suggested that card sorting exercises can be an important tool *when starting new projects*, in a new situation, or with a new group of people. Others suggested that they may be similarly useful as a reference *when evaluating a project*. Five participants suggested that card sorting exercises were an effective means of brainstorming, in particular *to structure understanding of an issue*. The activity of physically laying out cards on a shared work surface was seen as a useful way to find order, and relate different aspects to each other. Participants in group settings further observed that the cards were useful *to stimulate debate*, and think through a problem together. One participant stated, “I immediately started talking, they make you want to communicate, talk through the issue.” Another observed, “it’s easier to have the conversation with tools.”

Community engagement technique

- 16. Emphasise people's individual ability to make unique contributions.
 - 18. Address people individually instead of broadcasting your requests.
 - 42. Recognise different kinds of contributions to emphasise that there are many ways to participate.
 - 44. Give everyone an opportunity to speak, and prevent individuals from dominating meetings and activities.
 - 47. Foster social contact between contributors to increase enjoyment, and make tedious tasks more engaging.
 - 54. Nominate mentors: experienced participants who can guide newcomers.
-

Table 2. Frequently-chosen techniques.

4.3.1 Frequently Chosen Techniques. During the problem-solving sessions, some cards were chosen more frequently than others. Table 2 lists the cards that were most frequently chosen across multiple sessions.

4.3.2 Process Insights. The problem-solving sessions revealed three practical implications for the design and use of the CET catalogue.

Importance of expert facilitation. In some cases, techniques required context or clarification to ensure that they were understood and applied appropriately. When such facilitation was absent, we observed that participants interpreted cards in their own way, risking mis-interpretation.

Importance of a stated frame. Reviews of the full catalogue introduced many new ideas to a discussion. As a result, good framing was important (“what problem are we discussing?”), as was discipline in selecting actionable outcomes (“what interventions might address this specific problem?”). Some participants found that their understanding of the concern shifted as they reviewed its different aspects. For example, one participant found that the initially stated concern was an expression of an underlying deeper problem. In such cases it became important to regularly review whether the originally stated frame needed to be updated.

Design around loss aversion. Frequently, new ideas arose that were not specifically helpful for addressing the current concern, but were considered useful for later. This sometimes lead to loss aversion, where participants were hesitant to put a card aside when it did not apply to the imminent task. To address this, in later card sorting sessions we added a dedicated pile for such cards, providing a buffer for ideas that are worth considering at a later time.

5 SECOND VALIDATION STUDY

5.1 Evaluating Flickr Community Engagement

To complement the first qualitative study, in a second stage we tested the use for quantitative evaluation. Specifically, we sought to employ summative evaluation to identify group-specific practices that are associated with sustained contributions by new group members. This took the form of a large-scale observational study of 20,000 Flickr groups, with a total of 1 million members.

This second study had multiple objectives. We sought to determine whether CETs can be meaningfully translated to observable measures, and whether they could be used to describe group activities in a comprehensive manner. Most importantly, we sought to determine whether the catalogue of techniques can help interpretations of past engagement outcomes. This was intended as a demonstrator of a potential use of the catalogue, rather than a comprehensive evaluation of Flickr community engagement.

Category	Variable	Flickr indicator	CET #
Policy indicators	<i>privacy_public</i>	Is the group public?	5. Allow browsing, 13. Show activity, 39. Visible record.
	<i>privacy_invite_only</i>	Is the group invite-only?	25. Safe space.
	<i>privacy_private</i>	Is the group private?	25. Safe space.
	<i>show_pool</i>	Is the photo pool visible to non-members?	5. Allow browsing, 13. Show activity.
	<i>show_discussion</i>	Is the discussion forum visible to non-members?	5. Allow browsing, 13. Show activity, 47. Foster contact, 48. Varied encounter, 59. Communication tools.
	<i>has_description</i>	Does the group have a description?	21. Clear norms.
Group type	<i>num_admins</i>	Number of administrators and moderators	30. Grievance process, 57. Special powers, 54. Nominate mentors.
	<i>has_moderation</i>	Are photos reviewed before they are shared?	57. Special powers, 21. Clear norms.
Culture indicators	<i>groupref_per_member</i>	Comments with group reference to non-members	3. Outreach.
	<i>invites_per_member</i>	Invitations to non-members to join the group	3. Outreach.
	<i>threads_per_member</i>	Number of discussion threads per member	4. Regular updates, 47. Foster contact, 48. Varied encounter, 59. Comm. tools.
	<i>photos_per_member</i>	Number of photos shared per member	4. Regular updates.
Group type	<i>photos_gini</i>	Gini index of shared photos across members	44 Everyone contributes.
Group type	<i>st_score</i>	Raw social/topical group classification score	
	<i>st_social</i>	Social group? (<i>st_score</i> > 75 th percentile)	
	<i>st_topical</i>	Topical group? (<i>st_score</i> < 25 th percentile)	
	<i>st_mixed</i>	Mixed social-topical group? (All others)	
Control variables	<i>group_id</i>	Unique group identifier	
	<i>num_member</i>	Total number of members	
	<i>num_photos</i>	Total number of photos	

Table 3. Feature vector for the Flickr evaluation case study. Groups are characterised by their use of group policies, cultural factors that describe the nature of member contributions, and by their focus on social or topic-oriented activities.

5.2 Translation to Indicator Measures

To translate the techniques into indicator measures, the first and second author reviewed the CET catalogue, identifying aspects that matched the Flickr group context, and that could be derived from available Flickr metadata and contribution records. In total, the translation process took 1 hour. The two authors had prior familiarity with the domain: both are experienced participants in the community, and one author is a former employee and experienced researcher of the platform.

The derived set of evaluation indicators is described in Table 3. We distinguish three groups of techniques: policy indicators to characterise group rules, indicators of culture to characterise emergent behaviours of group participants, and systemic factors that hold for all Flickr groups. As the table shows, in some cases multiple CETs were coded with the same indicator variable.

Systemic factors are techniques shared by all Flickr groups, for example, capabilities for large-volume contributions (CET #27). These were not considered during evaluation.

11 CETs were coded in the form of *policy indicators*: these are engagement techniques that were enacted by group administrators, typically as group configuration options, and that affect all members of the respective group. This includes privacy settings to determine whether the group is public, private, or invite-only, and whether non-members can see the group's photo pool or discussion forum (CET #5,13,25,39). Other settings help manage the flow of contributions, and uphold social conduct. Administrators can require that all contributions are reviewed by a moderator before they are accepted (CET #21,57), and they can nominate other administrators and moderators to help manage the group (CET #30,54,57). The presence of a group description was interpreted as a means of providing guidance (CET #21.), and the presence of a group forum as a means to foster social interactions (CET #47,48,59).

6 CETs were coded as *cultural indicators*, reflecting engagement techniques that depend on emergent group behaviour. Several of these characterise overall group activity, including photo sharing (CET #4) and message volume in group discussions (CET #4,47,48,59). Two indicators

characterise collective efforts to promote the group to non-members (CET #3), either in the form of photo comments, or by making use of an explicit group invitation mechanism. One indicator characterises the relative distribution of contributions across members (CET #44), it was computed with a Gini coefficient.

5.3 Evidence Base

Our analysis encompasses the time period from Flickr's launch in 2004 until early 2016. To observe behavioural changes over time we defined three temporal study cohorts, based on an initial review of the Flickr contribution timeline. The cohorts reflect key periods in the Flickr history:

2007 cohort: an initial growth period from 1st Dec. 2006 - 1st Dec. 2007.

2010 cohort: a period of maturation from 1st Dec. 2009 - 1st Dec. 2010, this was later followed by a slow decline in photo sharing in 2011-2012.

2014 cohort: new sustained growth in photo sharing from 1st Dec. 2013 - 1st Dec. 2014, coinciding with a comprehensive redesign of the Flickr website in Summer 2013.²

In order to observe *engagement outcomes*, we built a statistical model of newcomer survival for every group. We identified instances where Flickr users joined a new group, and observed their activity in this group for a period of 360 days from the moment of joining, with a further observation period of 180 days to mark death events for right censoring. Newcomers are considered to be active in a group as long as they engage in at least one of three actions: sharing of photos in groups, commenting on the photos of other members, and marking such photos as favourites.

For statistical analyses, these group-specific survival models were then used to predict the likelihood of newcomer survival after 180 days of group membership. This latter 180-day threshold was chosen as a trade-off: it is long enough to measure sustained engagement over longer periods, but short enough that there still is some remaining activity. In preliminary analyses, survival distributions after 30, 90, and 180 days followed the same overall trend across the observed groups.

An observational study of Flickr participation has to consider a range of potential confounding factors, and the present study accounts for several such effects through control variables and group type classification. We account for 'rich-get-richer' phenomena in social networks [5], where larger groups tend to grow more quickly than others, by controlling for overall group size, group activity volume, and relative group age.

Differences in group purpose may lead to differences in emergent practice. In research of Flickr groups, an important distinction is made between groups with a social or topical focus [2]. Social groups show a higher degree of interactions and interconnectivity between members, while topical groups place an emphasis on content curation around particular topics. This difference in group focus may yield differences in emergent newcomer behaviour. To account for such effects, we classified groups according to their social and topical focus based on methods published in previous studies [2], and included the resulting features in our analysis.

We compiled a feature vector per group, consisting of all group-wise policy indicators, indicators of culture, control variables, group type features, and engagement outcome measures. We compared their relationships in a pairwise correlation analysis, and trained regression models to determine which groups were more likely than average to sustain newcomer activity. We further fitted a random forest regressor to assess relative feature importance.

²<http://blog.flickr.net/en/2013/05/20/a-better-brighter-flickr/>

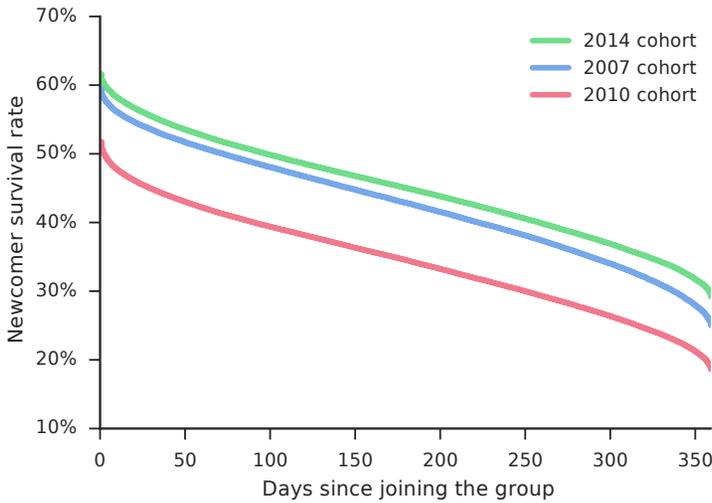


Fig. 4. Newcomer survival function per study cohort, from the moment of joining a new group. The plot visualises the share of newcomers who are still sharing photos, commenting, or favouriting photos after some time has passed.

5.4 Results

To establish a general frame for our evaluation, we first compared sustained newcomer engagement across the three temporal cohorts. The results of a survival analysis are shown in Figure 4. The 2007 and 2014 cohorts show the highest newcomer survival rates, while newcomers in the 2010 cohort could not be retained for as long. This suggests that newcomer retention rates in Flickr groups had decreased just before the 2013 redesign, but increased again after the redesign.

In a summative evaluation exercise, we then assessed to what extent this change in retention coincided with a change in behaviour. Figure 5 compares newcomer survival rates across the cohorts for three activities: photo-sharing, commenting, and favouriting. The plot shows a distinct change in contribution behaviour over time: sustained contribution activity by the early 2007 cohort more focused on social interactions, in that commenting activity by newcomers was sustained for longer periods than for the later cohorts. By comparison, commenting and favouriting activity for the 2010 cohort was markedly reduced, coinciding with an overall reduction in newcomer engagement. In contrast to this, newcomer retention in the 2014 cohort was much improved, and activity more focused on sustained photo sharing and favouriting. These results suggest that the 2013 Flickr redesign may have introduced a distinct change in contributor behaviour. In terms of photo sharing and favouriting activity, the 2014 cohort represents the most active Flickr community to date. However compared to prior cohorts, newcomer activity is less focused on social interactions between group members.

Results of a pairwise correlation across the feature vector suggest that the community engagement techniques make a significant difference: most policy and culture indicators are positively associated with sustained engagement. There was one exception: groups marked as private or invite-only had lower newcomer survival rates ($\rho_S = -0.27$ and $\rho_S = 0.95$ for photo sharing activity, respectively). On the other hand, public visibility of the photo pool and discussion forum had no significant impact on newcomer contribution activity, likely because this difference only affects non-members.

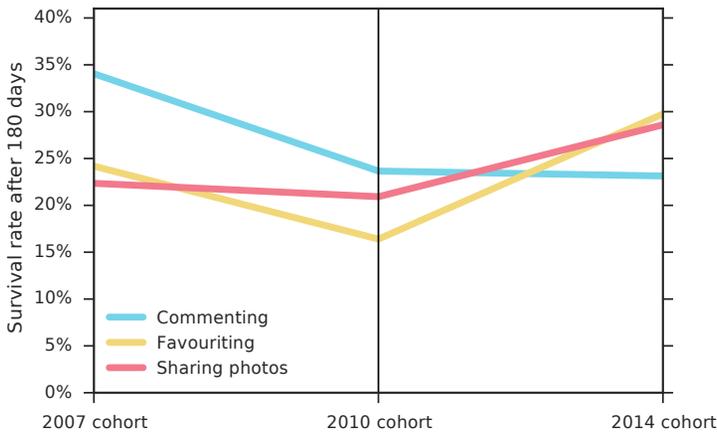


Fig. 5. Parallel coordinates plot: change in sustained newcomer behaviour across cohorts, for three types of contribution activity. The plot visualises the share of newcomers who still engage in an activity 180 days after joining a new group, as measured with a survival analysis.

Overall, correlation analysis showed that sustained newcomer engagement mostly takes place in very active groups. These are characterised with a larger member body ($\rho_S = 0.63$), a large number of photos in the pool ($\rho_S = 0.65$), and an unequal distribution of contributions, as measured with a Gini index across contributors ($\rho_S = 0.74$, all with $p < 0.001$). In social groups, members were slightly more likely to sustain commenting activity ($\rho_S = 0.11$) and favouriting ($\rho_S = 0.078$), while a comparison of their photo sharing activity was inconclusive. Groups that launched more recently pursued similar policies as older groups, however they were less likely to sustain photo sharing activity ($\rho_S = -0.15$), commenting ($\rho_S = -0.17$), and favouriting ($\rho_S = -0.11$, all with $p < 0.001$). This supports the observation that contributor behaviour has changed after the 2013 redesign.

To determine whether the pairwise correlations are indicators of real effects, we further assessed these relationships with an OLS regression model to explain sustained newcomer activity. Table 4 shows regression results for four models: an overall model across the full study population, and one model each for the three temporal cohorts. All regression models sought to explain the share of newcomers who are still active 180 days after joining a new group. The model fit was best for the overall model ($adjustedR^2 = 0.26$), and lower but comparable for the the cohort models ($adjustedR^s = 0.19, 0.17, 0.15$).

The regression results indicate a relative importance of culture indicators, rather than policy indicators, for improved engagement outcomes. They also provide further evidence for a shift in contributor behaviour over time, away from social interactions towards a focus on photo-sharing. The most dominant effect across cohorts relates to the distribution of photo-sharing activity, this was strongest for the 2014 cohort ($\beta_{photos_gini} = 0.44$), and less strong for the 2007 and 2010 cohorts ($\beta_{photos_gini} = 0.28, 0.34$). Conversely, the effect of outreach activity by group members was strongest for the 2007 cohort ($\beta_{groupref_per_member} = 0.05$). For the later cohorts, such outreach activity had no strong effect ($\beta_{groupref_per_member} = 0.00, 0.01$, all with $p < 0.001$).

We fitted a random forest regressor to predict which groups were more likely than average to sustain newcomer photo sharing activity. The resulting model allowed us to identify the features that most strongly predicted sustained activity. Table 5 lists the ten highest-ranking features in terms of importance, these account for 94% of total feature importance.

Category	Variable	Overall	2007	2010	2014
Policy indicators	<i>privacy_public</i>	-0.03**	-0.05*	-0.06**	(-0.03)
	<i>privacy_invite_only</i>	(-0.00)	-0.01	-0.04*	-0.08**
	<i>privacy_private</i>	-0.02**	-0.05*	-0.05*	-0.05*
	<i>show_pool</i>	0.05**	-0.04	0.06*	-0.04
	<i>show_discussion</i>	(-0.00)	-0.05	(-0.01)	(-0.01)
	<i>has_description</i>	0.01*	-0.03	-0.02	(-0.01)
	<i>num_admins</i>	(-0.00)	(-0.00)	(-0.00)	(0.00)
	<i>has_moderation</i>	(-0.00)	(-0.02)	(-0.00)	-0.04*
Culture indicators	<i>groupref_per_member</i>	0.00**	0.05**	0.00*	0.01*
	<i>invites_per_member</i>	(-0.00)	-0.03**	-0.03*	(-0.01)
	<i>threads_per_member</i>	0.01**	(-0.00)	0.02*	(0.00)
	<i>photos_per_member</i>	0.00**	0.00**	0.00**	0.00**
	<i>photos_gini</i>	0.28**	0.28**	0.34**	0.44**
Group type	<i>st_score</i>	0.02**	0.03**	0.03**	0.03**
	<i>st_social</i>	0.09**	0.09	0.12*	0.10*
	<i>st_topical</i>	0.07**	-0.07	0.08*	0.10*
	<i>st_mixed</i>	0.06**	0.08	0.11*	0.11**
Control variables	<i>group_id</i>	0.00*	(0.00)	0.00**	0.00**
	<i>num_members</i>	(-0.00)	(0.00)	(0.00)	(0.00)
	<i>num_photos</i>	0.00**	(0.00)	(0.00)	0.00*
	Intercept	-0.06**	-0.09*	-0.16**	-0.16**

Table 4. OLS regression results to explain sustained newcomer activity, measured as the share of newcomers who are still active 180 days after joining a group. This shows separate regression results for the full population, and the three temporal cohorts. Significant coefficients are either unmarked ($p < 0.05$), or marked with * ($p < 0.02$) or ** ($p < 0.001$). Coefficients in brackets are not statistically significant.

Variable	Type	ρ_S	Importance	σ
<i>photos_gini</i>	Culture	+	0.28	0.18
<i>num_photos</i>	Control	+	0.18	0.16
<i>photos_per_member</i>	Culture	+	0.13	0.06
<i>num_members</i>	Control	+	0.11	0.09
<i>st_score</i>	Group type	+	0.07	0.01
<i>group_id</i>	Control	+	0.07	0.00
<i>threads_per_member</i>	Culture	+	0.04	0.04
<i>privacy_public</i>	Policy	+	0.02	0.03
<i>invites_per_member</i>	Culture	+	0.02	0.00
<i>groupref_per_member</i>	Culture	+	0.01	0.00

Table 5. The ten highest-ranking feature importances and their standard deviations, as determined by a random forest regressor. The sign of the correlation coefficient ρ_S indicates the effect on newcomer engagement: a positive correlation indicates an increase in newcomer engagement when a particular feature is present.

5.4.1 Feature Importances. Many of the most important features are indicators of emergent group culture, such as the average contribution volume per member, the relative distribution of contributions among members, and the average number of discussion threads per member. Their relative prominence provides evidence that a ‘rich-get-richer’ dynamic is a fundamental aspect of sustained engagement: highly engaged groups are more likely to foster high newcomer engagement. Such an effect has been observed in other online settings [5, 11, 22, 31]. Its presence suggests that setting effective initial starting conditions can be key for subsequent sustained engagement.

Other important features are controls relating to group size and overall activity, providing further support for such an effect: larger groups and groups with more photo-sharing activity foster increased participation by new members. Only one policy indicator is among the most important ten features: the governance decision of making a group open to the wider public.

Overall, this suggests the presence of three important effects. First, it indicates an important trade-off between group privacy and sustained engagement. Second, it suggests that for the observed groups, collective and emergent group behaviour is more strongly indicative of newcomer engagement than most policy decisions made by group administrators. Highly engaged groups foster more engaged newcomers. Third, the observed participation inequality (where a small subset of contributors accounts for the largest share of contributions) is a common occurrence on volunteer platforms where contributions take place in an emergent manner, rather than being centrally coordinated [15].

5.4.2 Summary of Findings. Summative evaluation established that emergent cultural factors were more important for sustained newcomer engagement than policy decisions by group administrators. Newcomer retention was highest in larger and highly active groups that were open to the public. Groups with more sustainable newcomer engagement had set effective starting conditions for growth: their members engaged in practices that support newcomer recruitment, and the groups allowed for a long tail of participation. Taken together, this provided a more nuanced picture than the more general expectation of a ‘rich-get-richer’ phenomenon in social networks: it allows to distinguish the relative importance of different activities. Such findings can then inform future strategy, and help in the identification of development priorities.

The exercise also revealed that a platform redesign in 2013 effectively improved newcomer retention and contribution volume in Flickr groups, but this came at the expense of social interactions between group members. The reasons behind this surprising change in community behaviour were not immediately evident.

5.5 Development of Design Recommendations

We proceeded with a formative evaluation exercise to assess how commenting activity might be better supported. The aim of the exercise was to develop starting points for an intervention, primarily through the identification of potential root causes of the drop in commenting activity. It took the form of a combined card sort, identifying emergent thematic categories as we reviewed. The process took approximately 10 minutes. It yielded three actionable hypotheses about potential root causes:

- After the relaunch, photo comments were displayed less prominently in photo streams (CET #4,5,6,24,39).
- There currently is no simple way to keep track of ongoing discussions, making it hard to encounter commenting activity on other people’s photos (CET #7,13,27).
- More fundamentally, there currently is no explicit invitation to provide comments in the first place (CET #12,14,16,20).

As an unexpected outcome, the review also allowed us to develop a more nuanced understanding of the role of comments within the Flickr community. This helped us to better express how and why Flickr users might benefit from improved commenting interfaces:

- Comments can play an important role in peer support, for example to improve the learning experience of early-stage photographers (CET #32,33,49).
- Comments can help members bond over shared interests (CET #34,47,48,51,52).

Taken together, these observations provide rich starting points for interventions. Individual CETs yield specific design recommendations (for example, CET #13: “Show that others are already contributing”). Based on such findings, practitioners may choose to proceed with user studies to confirm certain hypotheses, or begin prototyping if there is already sufficient confidence in the benefit of a particular recommendation.

6 DISCUSSION

The first validation study was designed to assess suitability of the CET catalogue for applied problem-solving exercises. We sought to determine how long it takes to review the 59 techniques, whether the catalogue can provide suitable guidance, and whether it can support process innovation. The study also entailed open conversations about the potential use of the catalogue within existing organisational processes. We found that the catalogue can provide effective support for the design and review of community systems: a focused individual can review the 59 techniques in 15 minutes, the techniques act as effective prompts to reason about different aspects of a stated problem, and can help stimulate new ideas. When techniques are translated to method cards, the catalogue is highly suitable for in-person workshops, and it can assist conversations about complex issues in group settings.

The second study sought to assess use of the catalogue in quantitative evaluation. Can CETs be translated to observable measures, and can they help to interpret past engagement outcomes? For this study, we selected Flickr photo-sharing groups as a case study, and sought to identify policy choices and emergent cultural factors that have made an impact on sustained newcomer engagement. Using our method, we developed a more nuanced understanding of the impact of a large platform redesign, identifying both strengths and weaknesses of the new design. In a follow-up exercise, we provide an applied example of the use of our method to support strategic planning, developing broad design recommendations for future improvements. The outcomes illustrate the strength of our method: rather than offering mechanistic solutions to complex scenarios, it provides principled means of discovery and reflection, yielding nuanced observations that are grounded in both theory and practice.

In a more general sense, our evaluation demonstrated that the method presented here can provide effective support for the evaluation of community systems, and for triage of community engagement problems. The catalogue and card set can act as a detailed reference of considerations, and during review can help interpret past outcomes in a systematic and comprehensive manner. They can support both qualitative and quantitative evaluation, for example by translating CETs to indicator metrics that capture the use of particular engagement techniques.

Study outcomes also highlighted a need for appropriate contextualisation of the reviewed techniques. Organisers may interpret a technique according to their own understanding, which introduces a risk for misapplication. This is in part an artefact of the conceptual design of our catalogue: a CET describes ‘what has to be done’, but not ‘how to do it’. For this reason, appropriate application of CETs requires detailed knowledge of the observed context, and knowledge of the suitability of specific interventions. During our studies, an expert facilitator was present who was familiar with the material and could clarify terms and interpretations.

7 CONCLUSION

We believe that the method presented here makes a larger body of scientific knowledge more accessible to practitioners and researchers, and it can support lightweight reviews that are nevertheless systematic and evidence-based. The review process can apply to a wide range of scenarios, and can happen in a short amount of time: a single person can complete a full review in less than 15 minutes. During evaluations we found that the process yields effective prompts to consider particular concerns, that it yields a large number of plausible hypotheses, and that it can surface important and non-obvious findings.

A comprehensive catalogue of community engagement techniques is at the heart of the method, it is derived from existing collections in the literature. The catalogue is focused on immediate operational concerns of community organisations, including promotion and recruiting, provision of guidance for newcomers, solicitation and review of contributions, recognition and rewards, and shared social spaces. Based on the catalogue we develop a set of method cards, tangible artefacts which can be used in card sorting exercises for an individual or a group. For this purpose, we present simple exercises for formative and summative evaluation.

However, we also find that the current catalogue does not address many strategic and cultural concerns, including means to establish fundamental rules of interaction, negotiate priorities, select organisational structures, or determine organisational purpose. This is an artefact of its grounding in behavioural sciences, where studies focus on effects in individual and group psychology, rather than questions of organisational culture and purpose. In this sense, platform designers should consider the catalogue as one tool among many, and we encourage the development of further artefacts that support the design of community platforms in a more rigorous and systematic manner. To address these thematic limitations, future versions may broaden the scope of the catalogue to include insight from organisation science, management sciences, and other fields that study organisational processes. From a theoretical perspective, the CET catalogue could further become a starting point for more complex models of community engagement, including taxonomies of CETs in relation to particular concerns. There is a further opportunity to develop an empirical understanding of relationships across individual CET.

There review process itself presents additional design challenges that warrant further investigation. During our studies we relied on an expert facilitator who could clarify terms and interpretations, this helped avoid misapplications of techniques. Could the triage process itself incorporate a similar form of generalised guidance, for example by providing easy entry points into the literature, or by providing exemplary implementation strategies?

REFERENCES

- [1] Charles Abraham and Susan Michie. 2008. A taxonomy of behavior change techniques used in interventions. *Health psychology* 27, 3 (2008), 379.
- [2] Luca Maria Aiello. 2015. Group Types in Social Media. In *User Community Discovery*, Georgios Paliouras, Symeon Papadopoulos, Dimitrios Vogiatzis, and Yiannis Kompatsiaris (Eds.). Springer International Publishing, Chapter 5, 97–134.
- [3] Maria Antikainen, Marko Mäkipää, and Mikko Ahonen. 2010. Motivating and supporting collaboration in open innovation. *European Journal of Innovation Management* 13, 1 (2010), 100–119.
- [4] Jaime Arguello, Brian S Butler, Elisabeth Joyce, Robert Kraut, Kimberly S Ling, Carolyn Rosé, and Xiaoqing Wang. 2006. Talk to me: foundations for successful individual-group interactions in online communities. In *Proc. SIGCHI '06*. ACM, 959–968.
- [5] Albert-László Barabási and Réka Albert. 1999. Emergence of scaling in random networks. *science* 286, 5439 (1999), 509–512.
- [6] Jonathan Bishop. 2007. Increasing participation in online communities: A framework for human–computer interaction. *Computers in human behavior* 23, 4 (2007), 1881–1893.

- [7] Martin Boeker, Werner Vach, and Edith Motschall. 2013. Google Scholar as replacement for systematic literature searches. *BMC medical research methodology* 13, 1 (2013), 1.
- [8] Wichor M Bramer, Dean Giustini, Bianca MR Kramer, and PF Anderson. 2013. The comparative recall of Google Scholar versus PubMed in identical searches for biomedical systematic reviews: a review of searches used in systematic reviews. *Systematic reviews* 2, 1 (2013), 1.
- [9] E Gil Clary, Mark Snyder, Robert D Ridge, John Copeland, Arthur A Stukas, Julie Haugen, and Peter Miene. 1998. Understanding and assessing the motivations of volunteers: a functional approach. *Journal of personality and social psychology* 74, 6 (1998), 1516.
- [10] Dan Cosley, Dan Frankowski, Sara Kiesler, Loren Terveen, and John Riedl. 2005. How oversight improves member-maintained communities. In *Proc. SIGCHI '05*. ACM, 11–20.
- [11] Nicolas Ducheneaut, Nicholas Yee, Eric Nickell, and Robert J Moore. 2007. The life and death of online gaming communities: a look at guilds in world of warcraft. In *Proc. SIGCHI '07*. ACM, 839–848.
- [12] Andrea Forte and Cliff Lampe. 2013. Defining, understanding, and supporting open collaboration lessons from the literature. *American Behavioral Scientist* 57, 5 (2013), 535–547.
- [13] Gartner. 2012. Organizations That Integrate Communities Into Customer Support Can Realize Cost Reductions of Up to 50 Percent. (2012). <http://www.gartner.com/newsroom/id/1929014>.
- [14] Jean-François Gehanno, Laetitia Rollin, and Stefan Darmoni. 2013. Is the coverage of Google Scholar enough to be used alone for systematic reviews. *BMC medical informatics and decision making* 13, 1 (2013), 1.
- [15] Mordechai Muki Haklay. 2016. Why is participation inequality important? *European Handbook of Crowdsourced Geographic Information* (2016), 35–44.
- [16] Hahrie Han. 2014. *How organizations develop activists: Civic associations and leadership in the 21st century*. Oxford University Press, USA.
- [17] Caroline Haythornthwaite. 2007. Social networks and online community. *The Oxford handbook of Internet psychology* (2007), 121–137.
- [18] Quy Huy and Andrew Shipilov. 2015. The Key to Social Media Success Within Organizations. (2015). <http://sloanreview.mit.edu/article/the-key-to-social-media-success-within-organizations/>.
- [19] IDEO. 2003. method cards: 51 ways to inspire design. *Palo Alto* (2003).
- [20] Gerjo Kok, Nell H Gottlieb, Gjalt-Jorn Y Peters, Patricia Dolan Mullen, Guy S Parcel, Robert AC Ruiter, María E Fernández, Christine Markham, and L Kay Bartholomew. 2015. A taxonomy of behaviour change methods: an intervention mapping approach. *Health psychology review* (2015), 1–16.
- [21] Julia Kramer, Shanna R Daly, Seda Yilmaz, and Colleen M Seifert. 2014. A case-study analysis of design heuristics in an upper-level cross-disciplinary design course. (2014).
- [22] Robert Kraut, Sara Kiesler, Bonka Boneva, Jonathon Cummings, Vicki Helgeson, and Anne Crawford. 2002. Internet paradox revisited. *Journal of social issues* 58, 1 (2002), 49–74.
- [23] Robert E Kraut. 2003. Applying social psychological theory to the problems of group work. *HCI models, theories and frameworks: Toward a multidisciplinary science* (2003), 325–356.
- [24] Robert E Kraut, Paul Resnick, Sara Kiesler, Moira Burke, Yan Chen, Niki Kittur, Joseph Konstan, Yuqing Ren, and John Riedl. 2012. *Building successful online communities: Evidence-based social design*. MIT Press.
- [25] Karel Kreijns, Paul A Kirschner, and Wim Jochems. 2003. Identifying the pitfalls for social interaction in computer-supported collaborative learning environments: a review of the research. *Computers in human behavior* 19, 3 (2003), 335–353.
- [26] Cliff Lampe, Rick Wash, Alcides Velasquez, and Elif Ozkaya. 2010. Motivations to participate in online communities. In *Proc. SIGCHI '10*. ACM, 1927–1936.
- [27] Roz Lasker and Elisa Weiss. 2003. Broadening participation in community problem solving: a multidisciplinary model to support collaborative practice and research. *Journal of Urban Health* 80, 1 (2003).
- [28] Charlene Li. 2015. Why No One Uses the Corporate Social Network. (2015). <https://hbr.org/2015/04/why-no-one-uses-the-corporate-social-network>.
- [29] Chris Luebke. 2009. Drivers of change. *Appropriate Technology* 36, 4 (2009), 67.
- [30] Lynne C Manzo and Douglas D Perkins. 2006. Finding common ground: The importance of place attachment to community participation and planning. *Journal of planning literature* 20, 4 (2006), 335–350.
- [31] Cameron Marlow. 2004. Audience, structure and authority in the weblog community. In *International Communication Association Conference*, Vol. 27.
- [32] Susan Michie, Michelle Richardson, Marie Johnston, Charles Abraham, Jill Francis, Wendy Hardeman, Martin Eccles, James Cane, and Caroline Wood. 2013. The behavior change technique taxonomy (v1) of 93 hierarchically clustered techniques. *Annals of Behavioral Medicine* 46, 1 (2013).

- [33] Susan Michie, Maartje M van Stralen, and Robert West. 2011. The behaviour change wheel: a new method for characterising and designing behaviour change interventions. *Implementation Science* 6, 1 (2011), 1.
- [34] Michael Muller. 2012. Lurking as personal trait or situational disposition: lurking and contributing in enterprise social media. In *Proc. CSCW '12*. ACM.
- [35] Michael Muller, N Sadat Shami, David R Millen, and Jonathan Feinberg. 2010. We are all lurkers: consuming behaviors among authors and readers in an enterprise file-sharing service. In *Proc. GROUP '10*. ACM.
- [36] MyLife. 2013. Top 10 Failed Social Media Sites. (2013). <https://www.mylife.com/blog/uncategorized/failed-social-media-sites/>.
- [37] Oded Nov. 2007. What motivates Wikipedians? *Commun. ACM* 50, 11 (2007), 60–64.
- [38] Jennifer Preece and Ben Shneiderman. 2009. The reader-to-leader framework: Motivating technology-mediated social participation. *AIS Transactions on HCI* 1, 1 (2009), 13–32.
- [39] Paul G Schervish and John J Havens. 1997. Social participation and charitable giving: a multivariate analysis. *Voluntas: International Journal of Voluntary and Nonprofit Organizations* 8, 3 (1997), 235–260.
- [40] P Wesley Schultz, Stuart Oskamp, and Tina Mainieri. 1995. Who recycles and when? A review of personal and situational factors. *Journal of environmental psychology* 15, 2 (1995), 105–121.
- [41] John Wilson. 2000. Volunteering. *Annual review of sociology* (2000), 215–240.
- [42] Christiane Wölfel and Timothy Merritt. 2013. Method card design dimensions: a survey of card-based design tools. In *IFIP Conference on Human-Computer Interaction*. Springer, 479–486.
- [43] Kath Wright, Su Golder, and Rocio Rodriguez-Lopez. 2014. Citation searching: a systematic review case study of multiple risk behaviour interventions. *BMC medical research methodology* 14, 1 (2014), 1.

Received June 2017; revised July 2017; accepted August 2017