Check for updates

Anther: Cross-Pollinating Communities of Practice via Video Tutorials

Adam G. Emerson
adam.emerson@mavs.uta.edu
University of Texas at Arlington
Arlington, Texas, USA

Shreyosi Endow
shreyosi.endow@mavs.uta.edu
The University of Texas at Arlington
Arlington, Texas, USA

Cesar Torres
cearto@uta.edu
University of Texas at Arlington
Arlington, Texas, USA

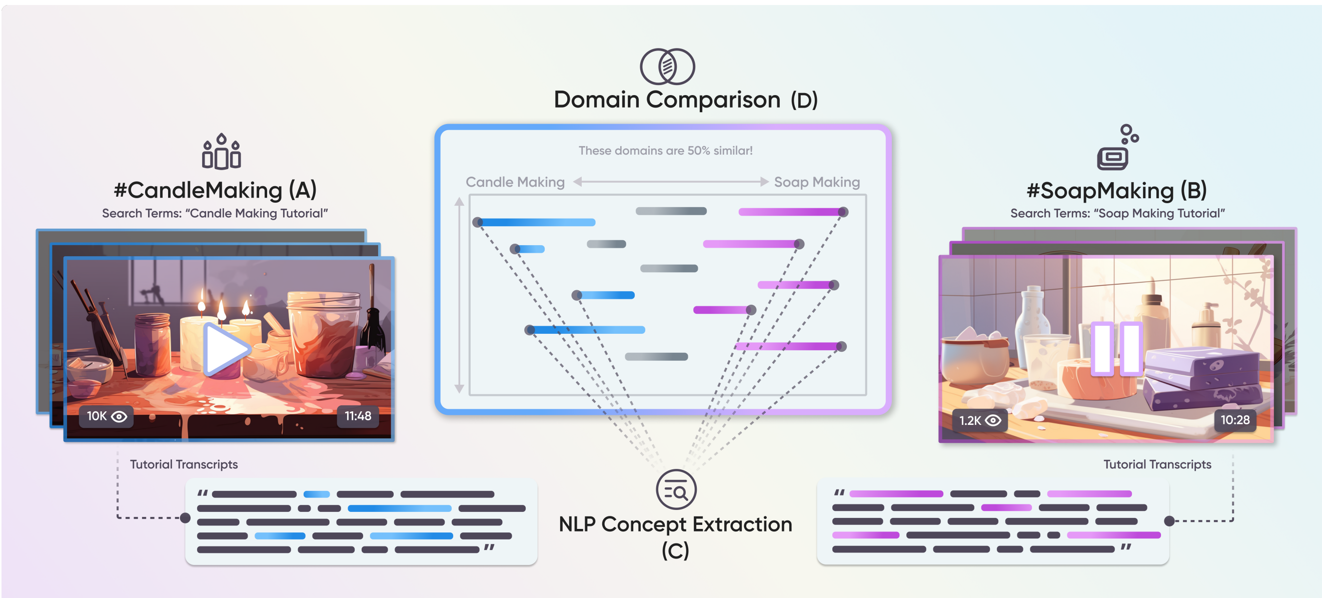

Figure 1: (A, B) Video tutorials are collected from candle and soap making communities-of-practice. (C) Domain concepts are extracted from tutorial transcripts using ML and NLP techniques. (D) The extracted concepts are visualized in a scatter-plot style tool, *Anther* for domain comparison and cross-pollination.

ABSTRACT

Communities of practice (CoPs) play a crucial role in cross-pollination and learning within various skill-based and craft domains. These communities often share common materials, concepts, and techniques across related practices. However, due to their insular nature, exchanging knowledge between CoPs has been challenging, leading to fragmented knowledge marked by differing vocabularies and contexts. To address this issue, we introduce Anther, a system designed to highlight shared concepts and semantic overlap between distinct CoPs. Anther projects concepts onto a 2-dimensional space, providing users with comprehensive, contextual, and conceptual views. We conducted a user study, demonstrating Anther's effectiveness in aggregating and disseminating community-based knowledge, bridging gaps between CoPs, and supporting the cross-pollination of knowledge between CoPs. Further, we present interaction vignettes that illustrate how Anther can ease entry into new domains and aide in discovering new creative techniques. This work can benefit maker communities by fostering collaborative knowledge-building across diverse domains.

CCS CONCEPTS

• **Human-centered computing** **Collaborative filtering**; • **Applied computing** *Fine arts*; • **Information systems** *Collaborative and social computing systems and tools*.

KEYWORDS

creativity support tools, cross-pollination, communities of practice, tutorials

ACM Reference Format:
Adam G. Emerson, Shreyosi Endow, and Cesar Torres. 2024. Anther: Cross-Pollinating Communities of Practice via Video Tutorials. In *Designing Interactive Systems Conference (DIS '24), July 01–05, 2024, IT University of Copenhagen, Denmark*. ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3643834.3660727

This work is licensed under a Creative Commons Attribution-Share Alike International 4.0 License.

DIS '24, July 01–05, 2024, IT University of Copenhagen, Denmark
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0583-0/24/07
https://doi.org/10.1145/3643834.3660727

1 INTRODUCTION

In the modern landscape of digital knowledge sharing, the ways in which communities of practice (CoPs) [37] share insights and experiences have seen significant transformation. Where personal blogs, forums, and static websites were once the primary sources of knowledge sharing and collaboration amongst DIY communities [54], video content has since taken over that role. YouTube is now the most commonly used DIY resource on the internet [35], a shift in modality that is perhaps unsurprising given the vast supply of content available on the platform – uploads have increased from 35 hours of video per minute in 2010 [60], to over 500 hours of video per minute in 2019 [25].

Video tutorials have been positioned as a form of informal "free-choice learning" [35, 49], where learning outcomes are guided by individual interest as opposed to a designated curricula or rubric. This model of learning parallels the ethos of CoPs, where membership is voluntary and the success of communities is driven by "their ability to generate enough excitement, relevance, and value to attract and engage members." [62]

CoPs have historically developed their own set of unique form-giving techniques inspired by "other adjacent and parallel practices, from which lessons are learned, innovations borrowed, procedures copied" [12, 61]. This process is sometimes referred to as "**cross-pollination**" – the exchange and integration of knowledge, techniques, and insights across CoPs with differing areas of expertise [23, 31, 64].

Cross-pollination creates opportunities for established practitioners within a CoP to expand their repertoire of form-giving techniques by borrowing from other CoPs and adapting them to fit their contexts. These forms of exchange can facilitate a deeper understanding of one's own practice, and sometimes, can even lead to new, tangential practices. E-Textiles, for example, could be conceptualized as a "seed" resulting from the cross-pollination of two parent practices - craft-based communities (e.g., knitting, sewing, embroidery) and engineering communities (e.g., electronics, computer engineering). The E-Textiles CoP demonstrates how the exchange of materials, tools, and techniques between seemingly disparate communities can engender a transformative practice with an identity that is distinct from its parent practices.

Yet, *digital platforms* are not always conducive to creative cross-pollination. Traditional information retrieval techniques like keyword searches [56] can be skewed by collaborative filtering algorithms that tailor content for users and inadvertently limit their exposure to new topics and communities [30]. Video segmentation [22] (chapter markers) can be useful in finding precise knowledge, but only if a maker generally knows what they are looking for. These current search and filtering practices pose challenges for practitioners who wish to branch out from familiar territories, limiting creativity and innovation.

In this work, we introduce a *Cross-Pollination Creativity Support Tool* (CP-CST) called Anther to explore how knowledge can be discovered and shared between communities by use of video tutorials. Traditionally, tutorials have acted as a form of *autogamy* or "self-pollination", primarily engaging existing members of a CoP. By uncovering and leveraging the shared semantics of CoPs, Anther extends the utility of tutorials beyond their original audience and facilitates *allogamy*, or "cross-pollination". Anther aims to promote growth and active participation within CoPs, but also to support the development of a more interconnected and collaborative making community at large. In contrast to approaches that provide explicit instruction to teach a new skill, Anther aims to aid learners in contextualizing an unfamiliar CoP by providing a bird's eye view of its most relevant concepts and identifying shared concepts with familiar CoPs.

Through our work, we contribute:

- *Cross-Pollination Creativity Support Tool* Anther, is an open source, WCAG-compliant[1], concept visualization and domain exploration tool designed to enable the combining of ideas, methods, and practices from different communities of practice. Our approach focuses on supporting tutorial reception [59], presenting learners with novel holistic, contextual, and depth-based views that aim to generate innovative solutions, creative breakthroughs, and new perspectives. Concepts are gathered from, and link back to video tutorials, offering a multi-dimensional learning experience. Concepts are visualized on a two-axis graph, where the primary axis is used to represent concepts relevant to a practice. For the secondary axis, we introduce a concept gravity score to indicate a concept's "pull", or prevalence throughout video tutorials.
- *Interaction Vignettes and User Study* We articulated the novel interaction techniques enabled by Anther using three interaction vignettes. We also conducted a user study with 10 participants (5 male, 5 female) from self-identified CoPs to better understand how Anther influences knowledge discovery and cross-pollination. We consolidate our findings into a set of descriptive themes that shed light on how Anther influences CoP navigation, user perception of cross-pollination strategies, and the way generative insights form and evolve through different conceptual views.

The paper first describes relevant research on tutorial-based community engagement and enhancement techniques, positioning Anther in relation to these methods as a novel approach for broadening the reach of video tutorials. We then provide three guiding design principles motivated from related work for the development of cross-pollination creativity support tools. We detail the design of Anther for viewing and analyzing 27 distinct online communities and provide a set of interaction vignettes to communicate its use. We then describe our user study design and present results as set of themes; we use these themes to discuss opportunities for further cross-pollination interactions.

2 RELATED WORK

We describe related work that investigates facilitating community cross-pollination and tutorial enhancement. We position our work alongside ongoing efforts to improve how learners navigate tutorials across different mediums and domains. We describe existing techniques that have been used to extract semantic and salient concepts from unstructured texts.

[1]WCAG compliance refers to the adherence to the Web Content Accessibility Guidelines (WCAG), a series of recommendations for making web content more accessible to people with disabilities.

2.1 Cross-Pollination Between CoPs

In the context of creativity, cross-pollination is defined as the sharing, adapting, and building upon the diverse expertise of others. Cross-pollination practice manifest across subjects and fields through networking, sound-boarding (i.e., good listening and feedback), collaboration, and win–win competition [33]. Exploring the dynamics of cross-pollination between CoPs offers insight into the challenges and opportunities for connecting making domains. Rakib et al. [47] demonstrated how co-designed ontologies had value in establishing a common language between glassworking and plastic 3D printing practices, yet forming these ontologies is costly, requiring diverse expertise and iteration. Chan et al. approached the cross-pollination problem through the lens of analogy-driven reasoning[6, 7, 16]. For example, SOLVENT [6] leveraged crowd-sourced annotations of research papers to identify analogies across distant domains and build semantic models to compare different fields. While annotations can offer rich insights into how users perceive information, sourcing these annotations for every new creative resource from the crowd remains costly. "Boundary objects" – or annotations that manifest as tangible and shareable artifacts – have been shown to be effective tools for CoP cross-pollination by stimulating conversation and creativity but these objects require careful curation and documentation to remain effective pollination tools [20]. Using video tutorials as a proxy to the common language of a CoP, Anther provides a computational method for extracting and linking common vocabularies and scaffolding entry experiences.

2.2 Mining Tutorial Content

Makers engage in dialogue not just to share information and encourage one another, but also to inspire, expecting that their ideas will be tweaked, remixed, and evolved [58]. Online tutorials often support community dialogue asynchronously through social media posts, comment sections, and metrics of endorsement ("upvotes" and "likes"). This content has been observed to include commentary and anecdotes on the hurdles, failures, and successes encountered while making, resulting in a conversational tone not seen in other forms of technical documentation [14]. Endow et al. [19] positioned these types of engagements as a form of patination – evidence of prior engagement and an indication of community sentiment. At their best, active comment sections elevate tutorials to something closer to a "living document", relying on community engagement to fill information gaps left by original authors and to keep material contemporary and accurate [32]. Some researchers have explored mining community comments and chat logs to improve tutorials [43]. Although comment sections can foster meaningful community dialogue, they can also lead to "toxic" behaviors including pedantry, bullying, and trolling [43]. In light of this, Anther does not perform analysis on video comment sections. Rather, we opt to link the users directly to video sources where comments can be explored on their respective platforms.

2.3 Tutorial Enhancement

Numerous works have aimed to improve the effectiveness of tutorials by better aligning them with the motivations of their target audiences. Some works propose frameworks for better tutorial authoring [12], while others propose new systems and taxonomies for improving tutorial reception [59, 63]. Lafreniere et al. [36] integrated a software directly into step-by-step tutorials, allowing users to interact with the program in question while following instructions. This system facilitated community engagement by recording users' actions and sharing them with the community for discussion, creating a feedback loop for tutorial improvement and evolution. ToolClips [24] supplemented existing application interfaces with video and text-based tooltips, augmenting users' understanding of tasks. DemoCut [9] applied machine learning techniques to reduce video tutorials down to their most essential parts, creating dense, information-rich videos with a quickly digestible format. Truong et al. converted existing makeup tutorials into hierarchical sets of instruction by applying principles of cognitive psychology, computer vision, and transcript analysis [57]. Pavel et al. [46] explored the generation of structured text from video content to provide an additional interface for video and tutorial navigation. Anther also relies heavily on transcript analysis, but rather than using it as a means to alter the tutorial structure, it uses the data to synthesize comprehensive overviews of domain knowledge. Our work aims to augment user's understanding of tasks by providing domain-specific information through a meta-layer of text-based insights and graphs.

2.4 Concept Extraction

Concept extraction from unstructured texts creates opportunities to reconfigure information in novel ways that improve user experiences and learning outcomes. Scatter/Gather [26] clusters documents within a corpora based on a given term, displaying shared "topical terms". ConceptScope [65] builds on this idea by leveraging ontologies, but these are static sources that require maintenance and intervention. Anther extract concepts from a dynamic source of documents (video tutorials) ensuring that the corpora remains contemporary to its domain. Working with video tutorial transcripts requires filtering out noise in order to extract valuable concepts. Researchers have employed probabilistic models to extract multi-word concepts or phrases. El-Kishky et al. [18] introduced the ToPMine algorithm which efficiently mined candidate phrases from corpora, outperforming other attempts at the same task. Li et al. [40] explored using vector embeddings to mine concepts using semantic context. Smith et al. [52] used crowd-sourced labels to evaluate probabilistic topic model visualizations. Chandrasegaran et al. [8] leveraged topic modeling to identify themes within group meeting discussions and computed the discussion "relatedness" to items in the meeting agenda using word embeddings. Scim [21] makes use of a language model to categorize and highlight documents at the sentence level for easier reading. While these probabilistic methods are effective ways to highlight and extract concepts, Anther leverages part-of-speech tagging, stop-word lists, and term-weighting as a low cost alternative. Term Frequency Inverse Document Frequency (TF-IDF) [53] is a popular term-weighting technique which highlights the words best suited to identifying a single document in a corpus. An alternative term-weighting scheme, Term Frequency Proportional Document Frequency (TF*PDF) [5], is often used as a tool for identifying emerging topics in a target domain [45]. Anther

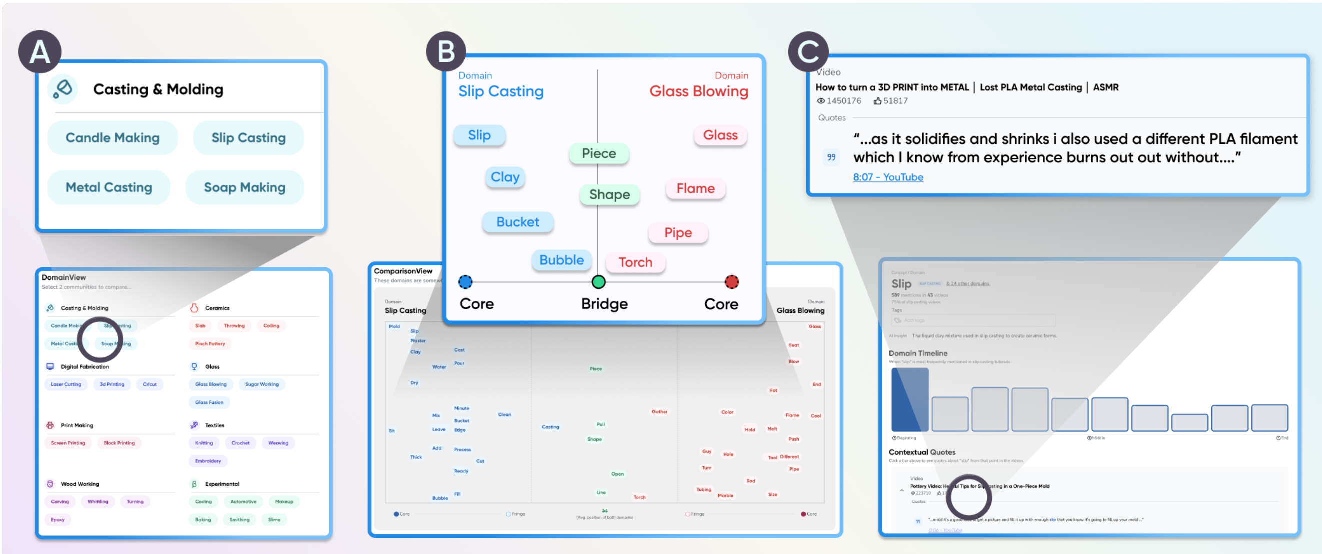

Figure 2: Anther Views: (A) DomainView groups maker communities into larger categories based on their linguistic correlations, (B) ComparisonView visualizes the conceptual correlations between two domains, and (C) ConceptView provides contextual overviews of how concepts are positioned within a domain and links back to source tutorials.

makes use of the TF*PDF algorithm to identify key concepts in each domain. TF*PDF is used alongside a derived metric, *Gravity*, to enable the visualization of extracted concepts in a two-dimensional space.

3 TOOL DESIGN PRINCIPLES

In order to develop a cross-pollination creativity support tool (CP-CST) for enabling concept exchange between CoPs, we identified three guiding design principles motivated from related work, each addressing a critical aspect of the cross-pollination process within tutorials. We use these guiding principles to inform the design of our CP-CST to better facilitate access to information and support the nuanced needs of individually motivated communities.

- **Guided Discovery** Conversation and dialogue are the foundation of any community, serving to establish a common language between community members and contributing to a sense of shared identity. Tutorials support this foundation in several ways – by providing clear documentation[14], allowing for direct and contextual feedback [44], and connecting members through authorship and viewership [59]. Tutorials have been understood as playing an important role in community dialogue by bringing in an informational meta-layer that accompanies a maker's artifact and facilitates discussion [54].
  CP-CSTs should enable users to navigate and connect with various CoPs. By providing cues that reveal interconnections between practices, we can assist makers in contextualizing new information against their prior experience. Such tools should provide the user with opportunities to engage with unfamiliar concepts without overwhelming them. While sparks of inspiration are critical for initiating creative processes, they are insufficient for achieving productive outcomes. CP-CSTs should not only ignite these sparks but also provide a "fuse" – a pathway of guided discovery, directing creative impulses towards tangible results.
- **Identity Bridging** Systemic biases, injustices, and cultural differences were shown to alienate makers and motivate the creation of separate making communities [50]. Even in carefully designed makerspaces, physical spaces meant to facilitate interdisciplinary making, knowledge remains distributed amongst makers and making 'zones', with true co-location of practices proving difficult to accomplish [1]. Furthermore, the feeling of 'otherness' within CoPs (e.g., imposter syndrome) and between CoPs (e.g., gatekeeping) influence participation in different practices [55]. These barriers disproportionately affect casual makers who are more easily discouraged upon encountering failures, and as a result, dissuaded from further exploration [28].
  CP-CSTs should champion inclusive practices, making sure that users of all abilities and backgrounds can benefit from using the tool. The tools should instill in users a sense of belonging and self-worth and mitigate alienation from perceived "otherness". Users should be empowered to inhabit the spaces between traditionally defined roles (e.g., "Artist-Engineer" rather than strictly "Artist" or "Engineer"), eliminating the need for code-switching as a means of validation and encouraging users to embrace multifaceted identities within their CoP.
- **Semantic Symmetries** Terminology and jargon act as barriers to entry, making even basic tutorials seem inscrutable to newcomers. This is compounded when CoP resources, like video tutorials, do not account for the varied backgrounds and skill sets of learners. These knowledge gaps have been

shown to result in siloed learning experiences, where learners hesitate to venture outside their comfort zones and instead stick to what they know [19].

CP-CSTs should highlight semantic symmetries between communities, helping users from one community understand the language or jargon of another. Recognizing that each community has its specialized language, CP-CSTs should affirm the value of the knowledge users bring with them, ensuring no one feels inferior due to unfamiliarity with specific terms. For instance, in the context of film production, what is referred to as a "C-47" is widely known outside the industry as a "clothespin". CP-CSTs should demystify these asymmetries, and make it clear that regardless of the name used, the value and function remain constant.

Building on these principles, we developed a CP-CST called Anther which aims to provide a foundation for effective and inclusive creative exploration and discovery, as well as foster collaboration across diverse communities of practice.

4 ANTHER

Anther is a user-facing web application designed to facilitate exploration and discovery across various communities of practice. It comprises several high-level views that enable users to explore the relationships between concepts within and across domains (Figure 2). These views serve as tools to cross-pollinate knowledge, technique, and understanding within and between domains. We detail each of these views in this section and note that access to Anther is available via a GitHub repository. [2]

4.1 Concept Extraction

In this study, we detail a systematic approach to data collection from various maker communities. This data serves as the foundation of Anther and its respective views.

Mining. Given a particular CoP (e.g., Knitting CoP), we gathered relevant video content using a uniform search term format "[CoP Name] tutorials" (e.g., "Knitting tutorials") from YouTube. Video transcripts and metadata were scraped using the unofficial Youtube-Transcript-API [13]. Video selection was initially limited to 100 videos per CoP, beyond which the relevance of content returned by the search notably decreased. Videos lacking substantial dialogue or consisting primarily of non-descriptive audio elements were excluded. The metadata collected included essential information such as video ID, title, view count, like count, and duration. Tutorial transcripts were obtained in JSON format, containing individual lines of text along with their start and end times.

Metrics. We lemmatized the transcripts and applied part-of-speech tagging using spaCy [27], extracting the nouns and verbs from each transcript. To identify the key concepts for each community, a *Relevance Score* was assigned to each word using the Term Frequency Proportional Document Frequency (TF*PDF) algorithm [5], (Eq. 1-2). [3] This approach, by favoring terms frequently mentioned across a wide array of documents, allows us to surface concepts that we use to indicate the foundational knowledge of each community of practice.

$$W_c = \sum_{d=1}^{D=d} |F_{c,d}| \cdot exp(\frac{n_{c,d}}{N_d}) \tag{1}$$

$$|F_c| = \frac{F_c}{\sqrt{\sum_{k=1}^{K=k} {F_k}^2}} \tag{2}$$

To understand the distribution and significance of concepts within video tutorials specific to one community, we introduced a *Gravity Score* alongside our relevance metric. While TF*PDF excels in identifying key terms within a corpus, the Gravity Score quantifies the "temporal weight" of these concepts, indicating how frequently and at what stages they are mentioned throughout a video. By calculating the normalized position of each concept occurrence within a tutorial (Eq. 3) and aggregating these into a histogram with predetermined bins (Eq. 4, Fig. 3C), we gain insights into the potential distribution of a concept—whether it plays a central role throughout the entire process or is more localized to a particular step or procedure.

$$p_{c,d} = \frac{c_{start,d}}{d_{duration}} \tag{3}$$

$$g_{c,d} = \sum_{i=0}^{k} I\left(\frac{B_{i,f}}{\max(B_f)} > \tau\right) \tag{4}$$

"Stitch" in Knitting, for example, has a high gravity score, indicating it is relevant from start to end in many tutorials. On the other hand, "Cast", referring to the act of "casting on" or creating the first row of stitches, has a low gravity score, being typically only mentioned towards the beginning of tutorials. A high gravity score suggests a concept is a **Core** aspect of the domain, consistently mentioned across the tutorial, whereas a low score identifies it as a **Fringe** concept, mentioned in more isolated segments. By employing Relevance and Gravity scores in tandem, we can project concepts into a two-dimensional space. These metrics serve as the foundation for our tool, allowing us to visualize both the prevalence of concepts in a community as well as their distribution throughout tutorials.

Our initial dataset was collected between July and September of 2023. The system is designed so that TF*PDF scores are calculated at the database level in a view – this means that as new videos are added to the system, rankings and scores will adjust in real time. To assess how the input of new videos may change the ranking of concepts, we added videos to the Metal Casting domain in May of 2024 and recorded the results. Videos were added using our original approach. The same search terms were used and we took the first 100 videos returned by the YouTube algorithm. Of these 100 videos, 38 were new to our dataset. We compared the rank and TF*PDF scores of the top 100 concepts in the Metal Casting domain before and after the new videos were included in the dataset and found the following:

- 9 out of 10 of the top ten concepts remained the same, although their positions varied slightly. The top 3 concepts exhibited no change in position.

[2]https://github.com/The-Hybrid-Atelier/Anther

[3]TF*PDF is often deployed across multiple channels or document sources, and so the frequency across each channel is normalized (equation 2). In our case, we only sourced from YouTube and so normalization was not necessary.

Concept	Before	After	Rank Diff.
metal	926.59	1317.07	-
cast	704.24	1066.64	-
mold	549.83	968.62	-
sand	386.06	724.33	+1
pour	413.39	591.01	-1
casting	316.50	476.87	+2
piece	327.70	441.98	-1
melt	318.05	422.66	-1
place	215.98	341.31	+1
silver	168.89	293.21	+6
cut	223.03	287.15	-2
clay	140.11	282.83	+13
hot	200.22	272.05	-2
clean	175.25	256.26	+1
ring	160.72	243.33	+2

Table 1: The top 15 concepts in Metal Casting after the addition of 38 new videos to the dataset, where "Before" and "After" are TF*PDF scores rounded to the nearest hundredth, and "Rank Diff." is the change in position

- 6 concepts fell out of the Top 100 (foundry, fit, rubber, hour, vacuum, block) and 6 new concepts were introduced (stone, important, channel, tool, box, torch).
- The largest gain in rank was +39 (stone) and the largest loss in rank was -32 (fire)

The lack of movement at the top of the rankings seems to indicate that the system does well in identifying the most relevant concepts, while more dynamic shifts in position towards the lower end of the rankings indicate that a larger dataset may be required to further solidify results.

4.2 Views

To organize the data, Anther was split into three hierarchical views: *DomainView*, *ComparisonView*, and *ConceptView*. To enable guided discovery, the views are structured so that users can first get a sense of the range of possibilities before steadily drilling down into the granular details of a CoP or concept as they follow their interests.

4.2.1 DomainView. We ran our concept mining algorithm on 27 CoPs (Table 2) across 2.7k video tutorials. We chose CoPs that make significant use of the video tutorial medium and spanned a diverse set of interests.

The **DomainView** is the first view in the Anther's flow. Each CoP is grouped into 8 discrete categories according to the similarity in their concepts (Figure 2A), providing an intuitive way for users to begin their exploration. Most categories are made up of domains that make use of similar materials (i.e., "Woodworking", "Ceramics", and "Textiles"). Other categories, like "Casting & Molding" encompass communities that make use of similar techniques. The DomainView's organization of communities serves as a logical starting point for users who wish to compare and contrast CoPs and is the first opportunity for inspiration to emerge.

Category	**CoP**
Casting & Molding	Candle Making
	Metal Casting
	Slip Casting
	Soap Making
Ceramics	Coiling
	Pinch Pottery
	Slab
	Throwing
Digital Fabrication	3D Printing
	Coding
	Laser Cutting
Glass	Glass Blowing
	Glass Fusion
	Sugar Working
Print Making	Block Printing
	Screen Printing
Textiles	Crochet
	Embroidery
	Knitting
	Weaving
Woodworking	Carving
	Epoxy Casting
	Turning
	Whittling
Experimental [4]	Automotive
	Baking
	Makeup

Table 2: A breakdown of the 27 CoPs and their respective categories as seen in the *DomainView.*

4.2.2 ComparisonView. The **ComparisonView** (Figure 2B) provides users with a comprehensive visualization of the semantic similarities between two selected domains, denoted here as Domain 1 (d_1) and Domain 2 (d_2). The view consists of 3 regions that together make a single graph:

- A region on the left-hand side for concepts exclusive to d_1 - $\{c \in d_1 | c \notin d_2\}$
- A region on the right-hand side for concepts exclusive to d_2 - $\{c \in d_2 | c \notin d_1\}$
- A region in the middle for concepts that appear in both d_1 and d_2 - $\{d_1 \cap d_2\}$

The Concept Valley. The Top 30 TF*PDF ranked concepts from d_1 and d_2 (60 concepts total) are graphed according to their relevance and gravity scores within their respective domains. Concepts are arranged so that looking at the top corners will indicate the most relevant and core concept for each domain. Concepts in d_1 are color-coded blue and concepts in d_2 are color-coded red. The graph was restricted to a maximum of 60 concepts so that a sufficient overlap

[4] Several communities were included to demonstrate the generalizability of Anther, but did not cleanly fit into a parent category and were grouped together under the category "Experimental".

between domains could be demonstrated without overwhelming users and sacrificing the visual clarity of the graph.

The view is organized in such a way that the path of domain-specific concepts resembles a valley (∪). If you follow this path, descending from one "head" of the valley towards the center of the graph, you begin to encounter concepts less relevant and more fringe.

Bridge Words. Concepts occurring in both d_1 and d_2 have their relevance and gravity scores from each domain averaged. The result is a cluster of concepts in the center of the graph representing the linguistic similarity between domains. We use this to further reinforce our visual metaphor – a "Bridge" spanning our valley and connecting the domains on either side. Concepts in this region are color-coded green, acting as a signpost to concepts that may offer the path of least resistance when entering a new domain.

4.2.3 ConceptView. The **ConceptView** provides contextual and temporal overviews of how a given concept from the ComparisonView relates to a specific community (Figure 2C).

Concept Prevelance. At the top of the ConceptView, quantitative statistics tell a story on the usage and prevalence of a concept in the domain. Information on how many other communities mention the concept is also available, supporting identity bridging and exploratory interactions for inspecting how a concept is used in another community .

AI Insights. In order to provide additional context and disambiguate concepts, we leveraged an AI agent built on OpenAI's GPT-3.5-Turbo model to provide additional information to users on how each concept relates to its domain. The agent was provided with the concept and CoP name, and was asked to provide a clarifying short sentence (Table 3) which was then displayed to users.

Temporal Concept Density. In addition to statistics, temporal insights are presented, establishing *when* the concept is used by a community. We re-use the histograms generated to calculate a concept's gravity score (Eq. 4) to indicate the weight of each concept across the length of a video; we use this as a proxy to provide information about the characteristics of a practice's workflow. By examining this timeline, a concept can be viewed as *core*, or fundamental to a domain (i.e., uniform histogram) or *fringe*, pertaining only to a specific use-case or process (i.e., noticeable spikes in usage).

Tutorial Backlinks. Beyond temporal and quantitative metrics, the ConceptView includes a library of *Concept Quotes* where each reference to a concept in the source tutorials is highlighted as an excerpt. These quotes give users direct access to the knowledge shared by tutorial authors by providing a timestamped link to the source tutorial. Quotes are grouped by their relative position on the timeline of their source tutorials. The source tutorials are ordered by number of views to highlight the most credible videos that explore the selected concept. If a user finds a particular video to be irrelevant to their exploration, the video can be hidden so that quotes from that tutorial are no longer seen.

ConceptView offers an interactive exploration that enables users to delve deeper into specific concepts of interest and gain a richer understanding of their significance in their respective domains.

AI Insight Prompt - GPT-3.5-Turbo

You are a craft-based information assistant. Your job is to provided users with one-sentence responses to prompts that consist of a contextual domain and a one word token.

Your responses should follow theses rules:

(1) Do not use the words 'token' or 'domain' in your response.
(2) Create concise responses that are no longer than one sentence.
(3) Do not repeat the context in your response.
(4) If you cannot make a connection between the domain and the token reply with just "N/A"
(5) Use language that is appropriate for a professional setting and may be found in a dictionary.

Here is an example prompt and response:

Prompt:

```
domain: candle making, token: pour
```

One-Sentence Response:

```
The act of carefully pouring melted wax into molds or
containers to create candles
```

Now the real prompt:

Prompt:

```
domain: <DOMAIN>, token: <TOKEN>
```

One-Sentence Response:

Table 3: The prompt provided to GPT-3.5-Turbo to generate the AI Insight where <DOMAIN> and <TOKEN> are the variables provided.

5 INTERACTION VIGNETTES

To gain insights into knowledge navigation interactions across communities of practice, we leverage evaluation by demonstration [38] as a way to articulate the novel interaction techniques enabled by Anther. Each technique is described within a vignette with concepts seen in Anther denoted with the styling Concept. Concepts are annotated with their respective gravity scores (GS=X) to indicate their temporal prevalence in community tutorials. Relevant quotes from videos within each community of practice are pulled directly from Anther's *ConceptView*, annotated as CoP#X. This approach allows us to draw with more detail the role and design implications for cross-pollination interactions.

5.1 Identifying Concept Anchors

Part of the challenge of branching outside one's immediate community of practice lies in the unfamiliarity with other domains. An anchor, in its traditional sense, secures a vessel, preventing it from being swept away by currents or winds. Similarly, when

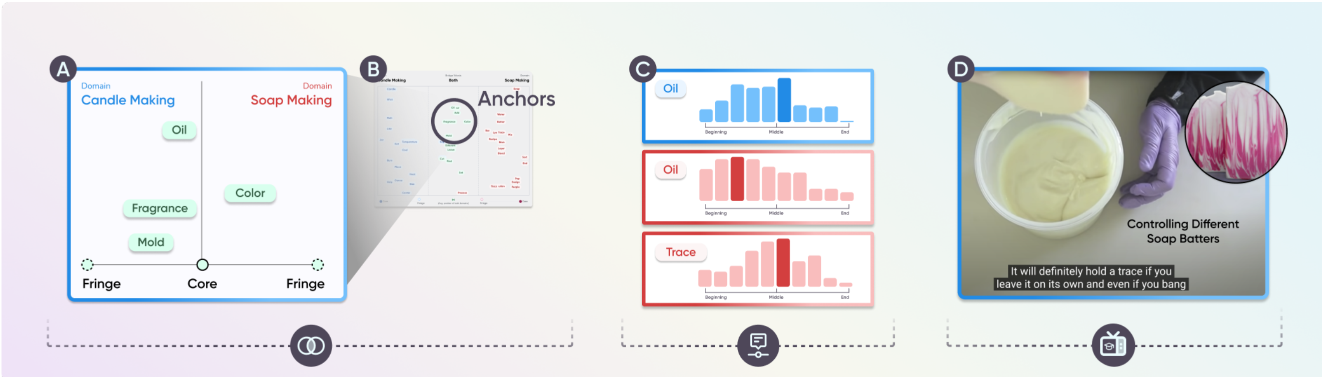

Figure 3: (A): A cluster of Anchor concepts positioned between the two domains. (B): A screenshot of the complete Comparison-View. (C): Concept timelines for Oil and Trace demonstrate variance in termporal relevancy. (D): A soap maker demonstrating their batter has reached the Trace state.

embarking on new intellectual or practical endeavors, users often seek "anchors" — familiar elements, concepts, or practices — that allow them to ground themselves within an unfamiliar domain. The *ComparisonView* highlights semantic symmetry, helping users to identify anchors by exploring intersecting concepts. For instance, consider a candlemaker aiming to venture into soapmaking. A cursory exploration reveals a triad of shared materials that can serve as anchors (Figure 3A-B): Oil (GS=7.5), Fragrance (GS=8.5), and Mold (GS=9).

Using these to ground the exploration, closer inspection can help disambiguate their analogous or divergent uses. Inspecting Oil and Fragrance materials more closely in the introduction of videos, both domains employ similar essential oils for fragrancing; however, moisturizing and lathering oils appear in new contexts:

> **SM1** There's a lot of oils in this. We're doing 11 ounces of **coconut oil**, nine ounces of **olive oil**, nine ounces of **palm oil**, two ounces of **sweet almond oil**, four ounces of **castor oil**. That is the sort of secret ingredient to get a lot of lather in your soap is to go pretty high on your **castor oil**. I'm really pushing the limits here. Four ounces is ten percent of this recipe, which is very high.[5]

Notably, the concept timeline indicates that Fragrance and Oil peak in usage mid-video (Figure 3C), suggesting its pivotal role during the casting phase for both crafts. A closer inspection reveals a new term:

> **SM2** . . . now the fragrance oil that I've chosen to use is a fragrance oil that doesn't accelerate **trace**. . . [6]

While the concept of Trace is mentioned across 21 CoPs, it assumes a special meaning in soapmaking:

> **AI Insight Trace** (*in soapmaking*) - The point in the soapmaking process when the mixture thickens and leaves a trail or "trace" when drizzled on the surface.

Unlike candlemaking where hot Wax is left to Cool, soapmakers work with mixtures that undergo a chemical reaction and thicken. The Thick and Trace concepts peak in the latter half of videos and reveal that the thickness of the solution is used to change the ways that materials are handled – viscous materials are spooned into molds, while thinner materials are poured, giving the soapmaker the ability to control how different soap batters interact with each other to give their characteristic swirl aesthetic (Figure 3D). While recipe-based exploration offers a prescriptive step-by-step guide to unfamiliar terrain, anchors provide a more flexible foundation, enabling individuals to navigate and understand new domains through familiar reference points.

5.2 Identifying and Tracking Whispers

Initially introduced as a technique, 3D Printing has since burgeoned into its own Community of Practice (CoP) and has been enriched by a plethora of interdisciplinary techniques[17]. However, these techniques are often eclipsed by content from the mainstream Maker community, resulting in hidden interdisciplinary conversations we term *whispers* .

Within Anther's ComparisonView, examining the 3D Printing CoP reveals a set of core concepts (GS=10), including Build, Model, Layer, and Filament. As the only material within the group, Filament appears as a fringe concept in a small number of other domains: Slip Casting, Metal Casting, and Laser Cutting (Table 4).

Concept	# Domains
Build	24
Layer	23
Model	18
Filament	3

Table 4: Core terms in 3D Printing domain and their correlation with other domains

This quality – ubiquity in one community and obscurity in all but a few others – is a signpost for emerging cross-pollination between CoPs.

Slip Casting offers only a single mention of Filament across 57 transcripts, appearing as a prominent spike on an otherwise empty domain timeline;

[5]https://www.youtube.com/watch?v=e4VU8onoc9E&t=123s
[6]https://www.youtube.com/watch?v=QKVSMfWn1m8&t=361s

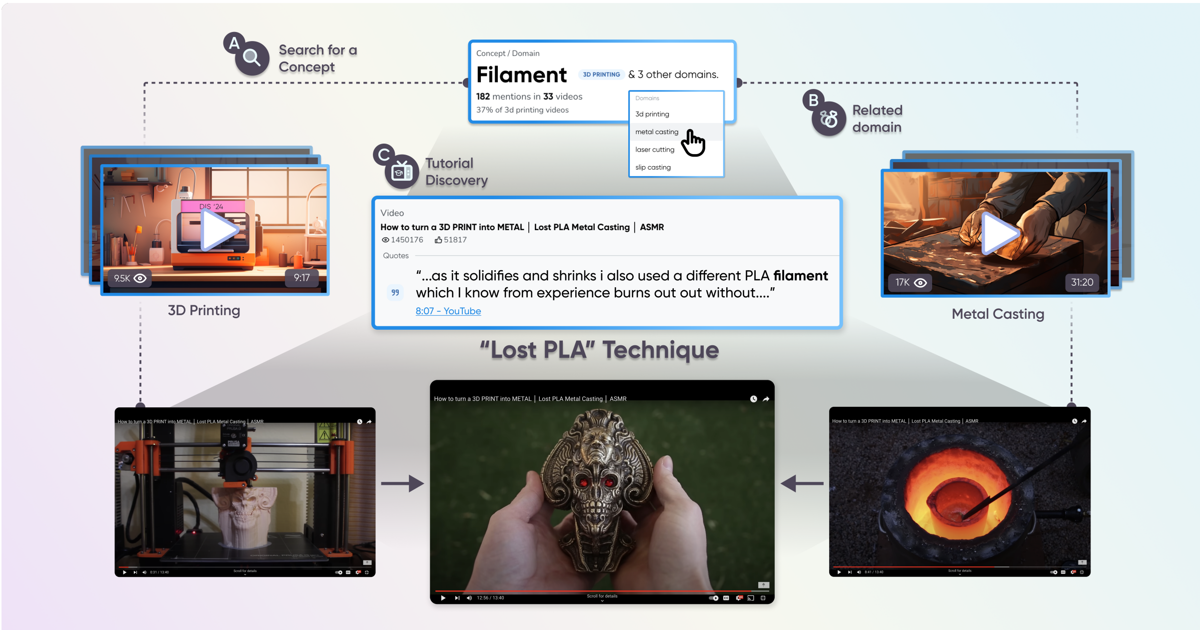

Figure 4: The process of discovering the "Lost PLA" technique, by Robinson Foundry, in which a PLA scuplture is coated in a ceramic mold and subsequently burned away with molten metal.

> **SC1** . . . you can use (cured) silicone to fill in the bulk, but it is an expense that just goes away versus **filaments** which is much cheaper or plaster which is really cheap . . .[7]

Here, we see the merits of casting different materials into 3D printed Molds. Silicone can be cost-prohibitive for makers and once set, cannot easily be repurposed. 3D printed molds, on the other hand, are much cheaper to produce, making them preferable in rapid prototyping or experimental contexts – this underscores the unique values and considerations within the slip-casting CoP when adopting 3D printing techniques.

Within Metal Casting, Filament is more frequently used. Several makers are experimenting with printing sculptures out of combustible Filament to create a Mold for their metal Casts;

> **MC1** I also used a different PLA filament which I know from experience burns out without leaving any ash behind. . .[8]

and

> **MC2** . . . and to do that the first step was to 3D print all of the puzzle pieces. Polymaker sent me this Polycast **filament** to try out. What's unique about it is that it can be burned away leaving virtually no ash behind.[9]

Both sculptors recognize the unique quality of PLA to burn away, only one refers to the technique as the "Lost PLA" method, a modern twist on the "Lost Wax" technique that has been in use since the bronze age. Anther thus allows access to both implicit and explicit cross-pollination by using materials, tools, or techniques to draw parallels between CoPs.

5.3 Seeding Creative Inquiry and Aesthetic Emulation

While many makers engage in crafting for the joy of the process itself, for most, the ultimate goal is to create a finished product. Aesthetic considerations, how an artifact looks, sounds, smells, and feels, often play a pivotal role in selecting the materials used for crafting. Materials that share little in terms of crafting technique can sometimes be substituted to achieve similar aesthetics. The choice of one material over another can be influenced by a variety of factors, including cost, availability, and requisite skill.

Anther facilitates the search for similar aesthetics through material exploration. By searching for concepts related to a desired aesthetic element, such as Glass, Anther reveals alternative domains where this aesthetic may be possible, like Glass Blowing, Sugar Working, or even Acrylic-based crafts. Aesthetic-driven material substitution is particularly evident in fields like glass blowing and sugar working. In these communities, several reasons establish the need for substitution: sugar is used instead of glass when the end product needs to be edible (cake decorations), when the end result is destined to be broken (prop making and SFX), or when glass artists need to practice with more readily available tools.

> **SW1** . . . if this is going to be kept around for a while, you guys don't want it to get cloudy. So . . . you can either use Dinky Doodle Shell and **Shine** spray, or you can use PME **Glaze** Spray.[10]

and

> **SW2** Well, that was exciting. Looked like **Glass**, sounded like glass, broke like glass. . . . I think the key here is to also use

[7] https://www.youtube.com/watch?v=ffj_1EqufE0&t=435s
[8] https://www.youtube.com/watch?v=rYNTua5fXxY&t=487s
[9] https://www.youtube.com/watch?v=UkjDwjaxkv8&t=25s
[10] https://www.youtube.com/watch?v=tfX6k-DQvig&t=1115s

metal on metal when you are stirring and just really keep it moving so you don't burn the **Sugar** and it doesn't turn that brown golden-yellow. [11]

Sometimes, aesthetic emulation requires clever adaptations. For instance, acrylic and glass are frequently substituted in contexts where the acrylic will be illuminated, such as in faux-stained glass and faux-neon artistry. Acrylic can mimic the appearance of glass tubes commonly used in neon signs. However, to achieve the glowing neon effect, these acrylic pieces require the addition of external factors like LED lighting. These supplementary elements are often offset by the cost or difficulty of implementing traditional methods.

These vignettes demonstrate novel techniques for knowledge navigation across CoPs that are made possible by Anther. To further evaluate how Anther influences creative exploration and knowledge discovery, we conducted a user study which we detail in the next section.

6 USER STUDY

We conducted a user study with 10 participants to assess how Anther influences cross-pollination between different communities of practice. We describe the study design and present a set of descriptive themes that emerged from analyzing interview and think-out-loud data.

6.1 Participants

Our pool of participants (10 total, 5 male, 5 female, Average Age: 22.2 ± 0.62 years) mostly comprised of college students who identify as hobbyists and practitioners for at least one craft practice. For the study, we defined hobbyists as those who are at the periphery of a CoP but are not fully integrated into the practice, and can therefore provide insights on entering a CoP. We recruited from our university's College of Engineering and College of Liberal Arts listservs and used snowball recruiting. A pre-screening questionnaire was used to gauge interest and determine participant's knowledge of CoPs supported by Anther (Table 2) by asking to self-report their knowledge on a scale of 0-100. Participants were recruited if they scored at least 30 or more on one of the provided practices. Our resulting participant pool demonstrated characteristics of individuals at the periphery of at least 6 of the 27 CoPs represented in Anther. The six CoPs included thrown pottery (ceramics), baking, makeup, 3D printing, candle making, and coding.

6.2 Study Design

Participants were asked to meet with us individually in-person in our lab. The study comprised of an hour long session where participants interacted with the Anther tool and were compensated in the form of a $10 Amazon gift card. We began the session by walking participants through the Anther tool and describing each of its features. Participants were then asked to explore the tool on their own, and examine different combinations of CoPs for around 5-10 minutes. During this process, participants were encouraged to follow a think-out-loud protocol and were subsequently audio and screen recorded. To avoid Hawthorne effects, no protocol personnel were present around the participant as they explored the tool. Once this exploration was complete, we presented participants with three tasks and simultaneously conducted semi-structured interviews:

- *Task 1* Participants were asked to choose two practices that they were familiar with. The interview questions probed on what they thought were similarities and differences between their chosen practices and if the tool reflected their thought processes.
- *Task 2* Participants were asked to choose one familiar practice and one unfamiliar practice that they are interested in. The interview questions probed on how they would generally get started with learning about the unfamiliar practice and how they envisioned Anther can complement the process.
- *Task 3* Participants were asked to choose two unfamiliar practices. The interview questions probed on what participants' experiences were in entering uncharted domains, the type of resources that are most beneficial, and how Anther can aid their navigation of unknown practices.

These tasks aimed to shed light on how Anther affected users' experiences of navigating familiar and unfamiliar CoPs.

6.3 Analysis

We transcribed audio recordings of participants' think-out-loud and interview sessions. From these transcripts, we extracted salient and insightful moments from which we generated memos. We conducted a thematic analysis on the memos leveraging Clarke and Braun's method [3, 11] where we iteratively developed and applied inductive codes to the memos. We then clustered the codes into a set of qualitative themes that shed light on how Anther concepts influenced CoP navigation, participants' perceptions of knowledge discovery strategies, and how generative insights formed and evolved through different concept views. We describe these themes below.

6.4 Themes

6.4.1 Theme 1: Materials and tools guide CoP navigation. Participants' perceptions of practices and how they vary were guided by the materials and tools that comprise the practices. Shared materials and tools influenced how likely participants were to pick up a new practice. One participant expressed how Anther was most useful when comparing sister practices (e.g. Slab Making and Pinch Pottery) within a larger domain (e.g. Ceramics) which aided in identifying shared tools and materials.

> **U1** If I am comparing embroidery and knitting, or throwing and slab, I would like to see the similarities because that way I know which supplies I have that are the same. It's easy to switch over to one or the other.

When exploring the ComparisonView, participants identified several concepts that described materials and tools for a given practice. Several participants expressed how an understanding of such concepts were fundamental to entering a new practice as a beginner. U2 described the lack of material and tool information as one of the barriers to starting a practice:

> **U2** You don't know what the right tool is. You don't know what material to use.

Some participants went beyond the bridge words presented in ComparisonView to identify tools shared by seemingly disparate

[11] https://youtu.be/WBxbvG4eEcU?si=pl1RB9Uejy17NkjS&t=262

practices. For example, comparing Candle Making and Coding, U4 said:

> **U4** I would think structure, or container, would be a common concept. Structure of candles would be like a jar or a mold, and in coding, we have data structures.

In some instances, we observed that concept exploration in ConceptView encouraged participants to consider how similar techniques can be modified for different materials. For example, U8, a cosplayer, investigated the concept stamp when exploring Block Printing, a domain unfamiliar to them. Upon examining the tutorials where stamp was mentioned, U8 discovered that the concept described the process of using a wooden stamp to transfer patterns onto fabric and said:

> **U8** Wait, this might actually be handy for cosplaying. When I did a cosplay, I kind of just painted the design onto the fabric and it was very tedious. But they are using wooden stamps, aren't they? I might need to go with a cheaper route though like a paper stencil. If I use a stencil, that would be a lot easier but I would need to be careful where I place it.

The insights reveal how materials and tools influence the mental models of practitioners and how material and tool literacy can spark cross-domain conversations.

6.4.2 Theme 2: Expertise and sensemaking ability shapes cross-pollination outcomes. Participants had mixed responses on where in their CoP navigation process they see the value of a CP-CST tool like Anther. Overall, participants agreed that compared to platforms such as YouTube, Anther got them closer to concept-relevant resources faster and more accurately.

> **U1** I think say you are looking at something super specific like embroidery, and you are having a big problem with a specific stitch, I think it's [Anther] useful because it's like a quicker way to get you to what you need without having to scroll through hundred of them.

Participants also highlighted how the presence of domain timelines and tutorial backlinks ensured that they got to relevant portions of tutorials without missing important concepts:

> **U2** Tutorials are often at a fast pace where they speed through the process and you end up missing a lot of points. This tool shortlists a lot of things so it gets to the right words and then to the right parts of a tutorial. It tells you if you are gonna use the concept in the beginning, middle or end, and I think it's better so you can prepare well and familiarize yourself with the process.

Some participants found value in using the Concept Valley in the Comparison View as a way to familiarize themselves with the most basic concepts within a practice and get a bird's eye view of the practice as a whole. U5 mentioned:

> **U5** If you are new to a domain, this tool gives you the most basic things you need to know for the domain. So it would be more useful in the beginning to get a larger glimpse of the domain.

Comparing the concepts to 'Most Replayed' graphs that visualize viewing activity on YouTube videos, U7 said:

> **U7** [In Anther] I will start with the most relevant concepts and go and look at the parts of the videos and kind of have an overall understanding. I would be able to just pick on the most important concepts instead of wasting my time about the details. I think it would be helpful because even without the system, I usually try to do that. So for example in YouTube videos, it shows which parts of the videos people viewed the most. So I go and watch them first.

However, some participants felt overwhelmed by the volume and generalized nature of concepts when exploring and learning a new practice. U1 suggested that practices in the DomainView should be split further to get more specific concepts:

> **U1** [On coding] I wish you could identify it by what language you use, like if you could further break it down to let's say an array in Matlab vs an array in R. You are already really overwhelmed in the beginning, so you don't want to know what's this concept across five different areas. I need to learn the skill I am looking for, and once I have learned it, I can explore a little bit.

U4 echoed the sentiment, suggesting that the concepts are valuable once they are familiar with the practice.

> **U4** In the beginning, if I know barely anything about crocheting, seeing a bunch of words like a lot of words connected to that subject might become a little overwhelming because I don't know where to start. That would be appreciated in the next step when I am branching out.

Interestingly, users had little to say about the AI generated insights. We speculate that this could be due to a lack of interactivity with the feature, or that user's attentions were already overwhelmed with the large amounts of data presented in the ConceptView.

These insights signal that the effectiveness of Anther as a CP-CST is influenced by sensemaking abilities and expertise. The features of the tool align with existing video tutorial platforms, while also exposing new ways of identifying and extracting information from tutorials.

6.4.3 Theme 3: CoPs share tacit similarities in their processes. When comparing CoPs, participants often identified similarities between processes within different practices that were not captured by the concepts. These similarities were often experiential and sometimes even metaphorical, focusing on the physical or mental act of being immersed in a process. For example, when comparing Coding and Baking, U7 mentioned how both practices have processes that value the aesthetic and visual qualities of the end product:

> **U7** I see similarities in the design process, like when you design a cake and you design the front end of a website or app, in both you want to make it beautiful and look good. Also sometimes, in both cases you are trying to impress a customer.

Similarly, U1 compared Throwing and Baking, highlighting the similarities between the physical act of wedging clay and kneading dough.

> **U1** If you are throwing, you can relate it to kneading dough. You kind of compact it to a ball and when you are trying to get the air bubbles out, you do the same thing in kneading.

On the other hand, participants also often identified differences between practices based on unspoken qualities such as their mental models and perspectives of the processes. Comparing Glass Blowing and Makeup, U2 described how shifts in perspective required for the processes make it difficult for techniques to cross-over:

> **U2** I don't think techniques can cross-over between these two domains. With glass blowing, you are looking at the glass as if in third person. Whereas when you are doing your own makeup, you are looking at a mirror so everything is inverted.

Such holistic insights about processes are difficult to capture with our concept extraction technique as they are often not verbalized by tutorial authors and therefore not externalized in tutorial transcripts. However, by exposing opportunities to compare CoPs, Anther aids in starting the conversation about potential similarities and dissimilarities between processes within a practice and inspiring deeper reflection of familiar processes.

7 DISCUSSION

We leverage our findings from the user study and interaction vignettes (referenced as *concept anchors*, *whispers*, and *aesthetic emulation*) to discuss future research trajectories for cross-pollination creativity support tools. We describe alternate trajectories for facilitating guided discovery through an unfamiliar CoP, discuss how richer social connections can aid in forming inclusive practices, and provide design implications for cross-pollinating with experiential and tacit concepts.

7.1 Directing Sparks

Findings from our user study signal how overlapping concepts between familiar and unfamiliar practices help spark an interest to explore a new practice. However, cross-pollination CSTs should go beyond the initial spark, exposing users to not only surface-level similarities, but also deeper reflections of practices. While Anther's ConceptView enables further investigation by linking users back to the original tutorials, participants cited other reasons preventing them from further exploration. The amount of time required to learn a new domain was raised as a hurdle, as well as the possibility that entering a new practice might distract them from their current one.

To support users in overcoming these challenges, personalized learning pathways within a CP-CST may be beneficial. For example, knowledge inventories can be integrated into the design of CP-CSTs to take stock of users' expertise and experiences and adjust accordingly. We see potential in leveraging Kaufman's 4C model of creativity [29] to identify the type of concepts that would be most suitable for users' in different development stages. Whereas a pro-C practitioner might benefit from contextual and niche concept overlaps with an unfamiliar practice, little-c non-experts may prefer generalized concepts that provide a birds' eye view of a practice to help them get started.

Similarly, material and tool inventories can aid in finding unfamiliar practices with a relatively lower entry overhead (*concept anchors*). Our study findings indicate that users may be more encouraged to dive deeper into a practice if they already possess the supplies required.

7.2 Championing Inclusive Practices

Anther enables users from all backgrounds and experiences to explore completely unfamiliar CoPs by highlighting commonalities with familiar domains. By shifting the focus on concept similarities, Anther does not make salient the cultural, social, and systemic nuances surrounding the CoPs. However, strong cultural separators still exist, as participants' perceive certain practices with pre-existing stigma.

As an example, although most of our participants compared Coding with craft based practices, they unanimously described coding as a practice with little room for creativity which is in stark contrast to research efforts in creative coding in HCI [2].

Future CP-CSTs like Anther can aid in clearing misconceptions and taking a step further in championing inclusive practices by sourcing concepts from a diverse body of instructional materials that capture non-traditional practices within a CoP. For example, within HCI, there is a significant overlap between textile based practices and programming [4, 15, 51], however, these whispers are not reflected on platforms such as Youtube without significant search term tuning (*whispers*). Participants mentioned using other platforms such as Reddit to guide them through a practice. Conversations in these forum-type platforms have potential for capturing social nuances associated with a practice. Additionally, this indicates a need for concept ranking strategies that visualizes niche but evolving concepts.

Furthermore, HCI researchers have described the value of physical presence when it comes to cross-pollination within CoPs [34]. Several participants indicated that when learning a new practice, they will often seek out a kindred soul to go through the process with them or an expert practitioner who can guide them through the practice. In Anther, we link the users to tutorial authors through tutorial backlinks. Future research trajectories can explore how richer social connections can be achieved by potentially connecting users to other learners in comments sections of tutorial videos.

7.3 Beyond Rudimentary Cross-Pollination

In this work, we leveraged part of speech tagging to extract concepts which were then ranked based on how frequently they appeared across several tutorials in the CoPs. An artifact of this approach was that the extracted concepts were often generic and signaled rudimentary similarities between practices, such as using one's hands in both practices. Although helpful in lowering the entry barrier to CoPs, insights from the user study suggest the need for concept extraction techniques that expose more contextual and experiential concepts that capture processes as a whole .

Additionally, future research trajectories should investigate concept extraction techniques that can identify anchors that are semantically dissimilar but behaviorally equivalent to further inspire users to find similarities between CoPs and probe deeper (*aesthetic emulation*). Such techniques can also have design implications for furthering existing research efforts in identifying material parallels which have been a driving force in expanding HCI's material horizons [47, 48]. Furthermore, we envision the scope of concepts to be expanded to capture the practitioners' environment which is fundamental in facilitating a dialog between a craftsperson and their practice [42]. Although information about the practitioners' environment is not always captured through textual transcripts, advances in augmented, virtual, and mixed reality applications can potentially be leveraged to encode spatial information which can later be translated into concepts [39, 41].

8 LIMITATIONS AND FUTURE WORK

Limited Study. Although our user study generated insights into the efficacy of Anther to support cross-pollination, we acknowledge that it is difficult to assess guided discovery in an one hour study. A longitudinal study is needed to better understand whether or not participants are empowered to dive deeper into new communities of practice.

Auto-Generated Transcripts. In the case that a tutorial author does not include a transcript with their video, YouTube will automatically generate one. While these transcripts are generally accurate, it's important to acknowledge they can be error-prone, particularly when trying to transcribe diverse accents. These inaccuracies highlight the need for supplementary verification methods to ensure the reliability of auto-transcribed data in Anther.

Search Term Tuning. For our data collection, we gathered tutorials using a uniform search term "[CoP] tutorials". While this produced relevant results, it is difficult to assess whether results obtained were optimal. Future work could be done to "tune" search terms to ensure that the language of each community is being used to gather their respective tutorials.

Participant Expertise. For this study, we examined how Anther may assist hobbyists and CoP newcomers enter a new creative practice. In the future, a second user study with CoP professionals can shed light on how Anther assists established community members with cross-pollination. Additionally, search patterns are known to differ with different age groups [10], and cross pollination behaviors may as well. Although our study comprised of college students, it is worth exploring how people in different age groups navigate Anther.

9 CONCLUSION

In this work, we presented Anther, a cross-pollination creativity support tool (CP-CST) that aides makers in exploring new and unfamiliar communities of practice (CoPs). Using video tutorials as a proxy for community dialog, we leveraged NLP techniques to extract key concepts for each CoP. We proposed design principles as a foundation for CP-CSTs and explored their value with Anther through illustrative vignettes. We evaluated Anther's efficacy as a CP-CST through a user study and consolidated our findings into a set of themes that demonstrate how Anther influences cross-pollination. A user study revealed that Anther enabled an exploration of CoPs that goes beyond rudimentary cross-pollination and encourages users to consider the materials, tools, techniques, and tacit processes that encompass both familiar and unfamiliar domains. Ultimately, we see Anther as an insightful probe into how future tools can empower collaborative making and cross-pollination between communities of practice.

ACKNOWLEDGMENTS

We would like to thank the reviewers for their insightful feedback and the users who made this research possible. This work was supported by the NSF REU Site on Hybrid Design and Fabrication (CNS-2150321).

REFERENCES

[1] Michelle Annett, Tovi Grossman, Daniel Wigdor, and George Fitzmaurice. 2019. Exploring and understanding the role of workshop environments in personal fabrication processes. *ACM Transactions on Computer-Human Interaction (TOCHI)* 26, 2 (2019), 1–43.

[2] Ilias Bergstrom and R Beau Lotto. 2015. Code Bending: A new creative coding practice. *Leonardo* 48, 1 (2015), 25–31.

[3] Virginia Braun and Victoria Clarke. 2021. Thematic analysis. Analysing qualitative data in psychology.

[4] Leah Buechley, Mike Eisenberg, Jaime Catchen, and Ali Crockett. 2008. The LilyPad Arduino: using computational textiles to investigate engagement, aesthetics, and diversity in computer science education. In *Proceedings of the SIGCHI conference on Human factors in computing systems.* 423–432.

[5] Khoo Khyou Bun and M. Ishizuka. 2001. Emerging Topic Tracking System. In *Proceedings Third International Workshop on Advanced Issues of E-Commerce and Web-Based Information Systems. WECWIS 2001.* 2–11. https://doi.org/10.1109/WECWIS.2001.933900

[6] Joel Chan, Joseph Chee Chang, Tom Hope, Dafna Shahaf, and Aniket Kittur. 2018. SOLVENT: A Mixed Initiative System for Finding Analogies between Research Papers. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 31 (nov 2018), 21 pages. https://doi.org/10.1145/3274300

[7] Joel Chan and David Anthony Rudd. 2023. Supporting Exploration of Far-Domain Analogical Inspirations with Bridging Analogies. In *Proceedings of the 15th Conference on Creativity and Cognition.* 293–297.

[8] Senthil Chandrasegaran, Chris Bryan, Hidekazu Shidara, Tung-Yen Chuang, and Kwan-Liu Ma. 2019. TalkTraces: Real-time capture and visualization of verbal content in meetings. In *Proceedings of the 2019 CHI conference on human factors in computing systems.* 1–14.

[9] Pei-Yu Chi, Joyce Liu, Jason Linder, Mira Dontcheva, Wilmot Li, and Bjoern Hartmann. 2013. DemoCut: Generating Concise Instructional Videos for Physical Demonstrations. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (St. Andrews, Scotland, United Kingdom) *(UIST '13)*. Association for Computing Machinery, New York, NY, USA, 141–150. https://doi.org/10.1145/2501988.2502052

[10] Jessie Chin and Wai-Tat Fu. 2010. Interactive effects of age and interface differences on search strategies and performance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems.* 403–412.

[11] Victoria Clarke, Virginia Braun, and Nikki Hayfield. 2015. Thematic analysis. *Qualitative psychology: A practical guide to research methods* 222, 2015 (2015), 248.

[12] Matthew A. Dalton, Audrey Desjardins, and Ron Wakkary. 2014. From DIY tutorials to DIY recipes. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems (CHI EA '14)*. Association for Computing Machinery, New York, NY, USA, 1405–1410. https://doi.org/10.1145/2559206.2581238

[13] Jonas Depoix. 2023. youtube-transcript-api. https://github.com/jdepoix/youtube-transcript-api.

[14] Derek Van Ittersum. 2014. Craft and Narrative in DIY Instructions. *Technical Communication Quarterly* (2014).

[15] Laura Devendorf and Chad Di Lauro. 2019. Adapting double weaving and yarn plying techniques for smart textiles applications. In *Proceedings of the Thirteenth International Conference on Tangible, Embedded, and Embodied Interaction.* 77–85.

[16] Zijian Ding, Arvind Srinivasan, Stephen MacNeil, and Joel Chan. 2023. Fluid transformers and creative analogies: Exploring large language models' capacity for augmenting cross-domain analogical creativity. In *Proceedings of the 15th Conference on Creativity and Cognition.* 489–505.

[17] Dale Dougherty. 2012. The maker movement. *Innovations: Technology, governance, globalization* 7, 3 (2012), 11–14.

[18] Ahmed El-Kishky, Yanglei Song, Chi Wang, Clare R. Voss, and Jiawei Han. 2014. Scalable Topical Phrase Mining from Text Corpora. *Proc. VLDB Endow.* 8, 3 (nov 2014), 305–316. https://doi.org/10.14778/2735508.2735519

[19] Shreyosi Endow and Cesar Torres. 2021. "I'm Better Off on My Own": Understanding How a Tutorial's Medium Affects Physical Skill Development. In *Designing Interactive Systems Conference 2021* (Virtual Event, USA) *(DIS '21)*. Association for Computing Machinery, New York, NY, USA, 1313–1323. https://doi.org/10.1145/3461778.3462066

[20] Gerhard Fischer, Elisa Giaccardi, Hal Eden, Masanori Sugimoto, and Yunwen Ye. 2005. Beyond binary choices: Integrating individual and social creativity. *Int. J. Hum. Comput. Stud.* 63, 4-5 (Oct. 2005), 482–512.

[21] Raymond Fok, Hita Kambhamettu, Luca Soldaini, Jonathan Bragg, Kyle Lo, Marti Hearst, Andrew Head, and Daniel S Weld. 2023. Scim: Intelligent skimming support for scientific papers. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. 476–490.

[22] C. Ailie Fraser, C. Ailie Fraser, Joy O. Kim, Joy O. Kim, Hijung Valentina Shin, Hijung Valentina Shin, Jonathan Brandt, Joel Brandt, Mira Dontcheva, and Mira Dontcheva. 2020. Temporal Segmentation of Creative Live Streams. *International Conference on Human Factors in Computing Systems* (2020). https://doi.org/10.1145/3313831.3376437

[23] STINE Grodal and GRID Thoma. 2008. Cross-pollination in science and technology: concept mobility in the nanobiotechnology field. In *NBER Conference on Emerging Industries: Nanotechnology and NanoIndicators*. 1–2.

[24] Tovi Grossman and George Fitzmaurice. 2010. ToolClips: An Investigation of Contextual Video Assistance for Functionality Understanding. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, Georgia, USA) *(CHI '10)*. Association for Computing Machinery, New York, NY, USA, 1515–1524. https://doi.org/10.1145/1753326.1753552

[25] James Hale. 2019. More than 500 hours of content are now being uploaded to Youtube every minute. https://www.tubefilter.com/2019/05/07/number-hours-video-uploaded-to-youtube-per-minute/

[26] Marti A Hearst, David R Karger, and Jan O Pedersen. 1995. Scatter/gather as a tool for the navigation of retrieval results. In *Working Notes of the AAAI Fall Symposium on AI Applications in Knowledge Navigation and Retrieval, Cambridge, MA*.

[27] Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. spaCy: Industrial-strength Natural Language Processing in Python. (2020). https://doi.org/10.5281/zenodo.1212303

[28] Nathaniel Hudson, Celena Alcock, and Parmit K Chilana. 2016. Understanding Newcomers to 3D Printing: Motivations, Workflows, and Barriers of Casual Makers. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) *(CHI '16)*. Association for Computing Machinery, New York, NY, USA, 384–396.

[29] James C Kaufman and Ronald A Beghetto. 2009. Beyond big and little: The four c model of creativity. *Review of general psychology* 13, 1 (2009), 1–12.

[30] Duk-jin Kim, Duk Jin Kim, Wooyoung Lee, Wooyoung Lee, Dohyung Kim, Do Hyung Kim, Gwang-Yong Gim, and Gwangyong Gim. 2021. An Empirical Study on Filter Bubbles in the YouTube Comments Network: Using Social Network Analysis. *International Journal of Software Innovation* (2021). https://doi.org/10.4018/ijsi.2021070104

[31] Daejung Kim and Jieun Kim. 2021. Is innovation design-or technology-driven? Citation as a measure of innovation pollination. *World Patent Information* 64 (2021), 102010.

[32] Juho Kim, Phu Tran Nguyen, Sarah Weir, Philip J Guo, Robert C Miller, and Krzysztof Z Gajos. 2014. Crowdsourcing step-by-step information extraction to enhance existing how-to videos. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Toronto, Ontario, Canada) *(CHI '14)*. Association for Computing Machinery, New York, NY, USA, 4017–4026.

[33] Kyung Hee Kim. 2019. Demystifying creativity: What creativity isn't and is? *Roeper Review* 41, 2 (2019), 119–128.

[34] Tiffany Knearem, Xiying Wang, Jiaan Wan, and John M Carroll. 2019. Crafting in a community of practice: Resource sharing as key in supporting creativity. In *Proceedings of the 2019 on Creativity and Cognition*. 83–94.

[35] Sean Kross, Sean Kross, Eszter Hargittai, Eszter Hargittai, Elissa M. Redmiles, and Elissa M. Redmiles. 2021. Characterizing the Online Learning Landscape: What and How People Learn Online. *null* (2021). https://doi.org/10.1145/3449220

[36] Benjamin Lafreniere, Tovi Grossman, and George Fitzmaurice. 2013. Community Enhanced Tutorials: Improving Tutorials with Multiple Demonstrations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) *(CHI '13)*. Association for Computing Machinery, New York, NY, USA, 1779–1788. https://doi.org/10.1145/2470654.2466235

[37] Jean Lave and Etienne Wenger. 2004. Communities of practice. *Learning, meaning and identity* (2004).

[38] David Ledo, Steven Houben, Jo Vermeulen, Nicolai Marquardt, Lora Oehlberg, and Saul Greenberg. 2018. Evaluation strategies for HCI toolkit research. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–17.

[39] Jiannan Li, Maurício Sousa, Karthik Mahadevan, Bryan Wang, Paula Akemi Aoyagui, Nicole Yu, Angela Yang, Ravin Balakrishnan, Anthony Tang, and Tovi Grossman. 2023. Stargazer: An Interactive Camera Robot for Capturing How-To Videos Based on Subtle Instructor Cues. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–16.

[40] Keqian Li, Hanwen Zha, Yu Su, and Xifeng Yan. 2018. Concept Mining via Embedding. In *2018 IEEE International Conference on Data Mining (ICDM)*. 267–276. https://doi.org/10.1109/ICDM.2018.00042

[41] Karthik Mahadevan, Qian Zhou, George Fitzmaurice, Tovi Grossman, and Fraser Anderson. 2023. Tesseract: Querying Spatial Design Recordings by Manipulating Worlds in Miniature. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–16.

[42] Maarit Mäkelä and Bilge Merve Aktaş. 2022. In dialogue with the environment: The environment, creativity, materials and making. *Craft Research* 13, 1 (2022), 9–34.

[43] Ilya Musabirov, Paul Okopny, and Denis Bulygin. 2015. Analyzing chat logs in online games for tutorial improvement. In *Proceedings of the 2015 Annual Symposium on Computer-Human Interaction in Play*. 661–666.

[44] Cuong Nguyen and Feng Liu. 2015. Making software tutorial video responsive. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 1565–1568.

[45] Khanh-Ly Nguyen, Byung-Joo Shin, and Seong Joon Yoo. 2016. Hot topic detection and technology trend tracking for patents utilizing term frequency and proportional document frequency and semantic information. In *2016 International Conference on Big Data and Smart Computing (BigComp)*. 223–230. https://doi.org/10.1109/BIGCOMP.2016.7425917

[46] Amy Pavel. 2019. *Navigating Video Using Structured Text*. PhD dissertation. University of California, Berkeley, Berkeley, CA. https://www.proquest.com/openview/662e08576bd294eb287c1959eacf50e8/1?pq-origsite=gscholar&cbl=18750&diss=y

[47] Mohammad Abu Nasir Rakib, Jeremy Scidmore, Justin Ginsberg, and Cesar Torres. 2023. Thermoplastic Kilnforms: Extending Glass Kilnforming Techniques to Thermoplastic Materials using Ontology-Driven Design. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 263–281.

[48] Erica Robles and Mikael Wiberg. 2010. Texturing the" material turn" in interaction design. In *Proceedings of the fourth international conference on Tangible, embedded, and embodied interaction*. 137–144.

[49] Sonny Rosenthal. 2018. Motivations to seek science videos on YouTube: free-choice learning in a connected society. *International Journal of Science Education, Part B* 8, 1 (Jan. 2018), 22–39.

[50] Daniela K Rosner and Sarah E Fox. 2016. Legacies of craft and the centrality of failure in a mother-operated hackerspace. *New Media & Society* 18, 4 (April 2016), 558–580.

[51] Daniela K Rosner, Samantha Shorey, Brock R Craft, and Helen Remick. 2018. Making core memory: Design inquiry into gendered legacies of engineering and craftwork. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–13.

[52] Alison Smith, Tak Yeon Lee, Forough Poursabzi-Sangdeh, Jordan Boyd-Graber, Niklas Elmqvist, and Leah Findlater. 2017. Evaluating visual representations for topic understanding and their effects on manually generated topic labels. *Transactions of the Association for Computational Linguistics* 5 (2017), 1–16.

[53] Karen Sparck Jones. 1988. *A Statistical Interpretation of Term Specificity and Its Application in Retrieval*. Taylor Graham Publishing, GBR, 132–142.

[54] Stacey Kuznetsov and Eric Paulos. 2010. Rise of the expert amateur: DIY projects, communities, and cultures. (2010).

[55] Cesar Torres, Sarah Sterman, Molly Nicholas, Richard Lin, Eric Pai, and Eric Paulos. 2018. Guardians of Practice: A Contextual Inquiry of Failure-Mitigation Strategies within Creative Practices. In *Proceedings of the 2018 Designing Interactive Systems Conference* (Hong Kong, China) *(DIS '18)*. Association for Computing Machinery, New York, NY, USA, 1259–1267. https://doi.org/10.1145/3196709.3196795

[56] Cristen Torrey, Elizabeth F. Churchill, and David W. McDonald. 2009. Learning How: The Search for Craft Knowledge on the Internet. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Boston, MA, USA) *(CHI '09)*. Association for Computing Machinery, New York, NY, USA, 1371–1380. https://doi.org/10.1145/1518701.1518908

[57] Anh Truong, Peggy Chi, David Salesin, Irfan Essa, and Maneesh Agrawala. 2021. Automatic Generation of Two-Level Hierarchical Tutorials from Instructional Makeup Videos. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 108, 16 pages. https://doi.org/10.1145/3411764.3445721

[58] Tiffany Tseng and Mitchel Resnick. 2014. Product versus Process: Representing and Appropriating DIY Projects Online. In *Proceedings of the 2014 Conference on Designing Interactive Systems* (Vancouver, BC, Canada) *(DIS '14)*. Association for Computing Machinery, New York, NY, USA, 425–428. https://doi.org/10.1145/2598510.2598540

[59] Ron Wakkary, Markus Lorenz Schilling, Matthew A. Dalton, Sabrina Hauser, Audrey Desjardins, Xiao Zhang, and Henry W.J. Lin. 2015. Tutorial Authorship and Hybrid Designers: The Joy (and Frustration) of DIY Tutorials. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) *(CHI '15)*. Association for Computing Machinery, New York, NY, USA, 609–618. https://doi.org/10.1145/2702123.2702550

[60] Hunter Walk. 2010. Great scott! over 35 hours of video uploaded every minute to YouTube. https://blog.youtube/news-and-events/great-scott-over-35-hours-of-video/

[61] Alan Warde. 2005. Consumption and theories of practice. *Journal of consumer culture* 5, 2 (2005), 131–153.

[62] Etienne Wenger, Richard McDermott, and William M. Snyder. 2002. *Cultivating communities of practice: A guide to managing knowledge*. Harvard Business School Press.

[63] Saelyne Yang, Sangkyung Kwak, Juhoon Lee, and Juho Kim. 2023. Beyond Instructions: A Taxonomy of Information Types in How-to Videos. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (<conf-loc>, <city>Hamburg</city>, <country>Germany</country>, </conf-loc>) *(CHI ’23, Article 797)*. Association for Computing Machinery, New York, NY, USA, 1–21.
[64] Anastasiya Yurchyshyna, Wanda Opprecht, and Michel Léonard. 2011. Collaborative decision constructing supported by Cross-Pollination Space. *Proc. COLLA* (2011).
[65] Xiaoyu Zhang, Senthil Chandrasegaran, and Kwan-Liu Ma. 2021. ConceptScope: Organizing and Visualizing Knowledge in Documents Based on Domain Ontology. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI ’21)*. Association for Computing Machinery, New York, NY, USA, Article 19, 13 pages. https://doi.org/10.1145/3411764.3445396