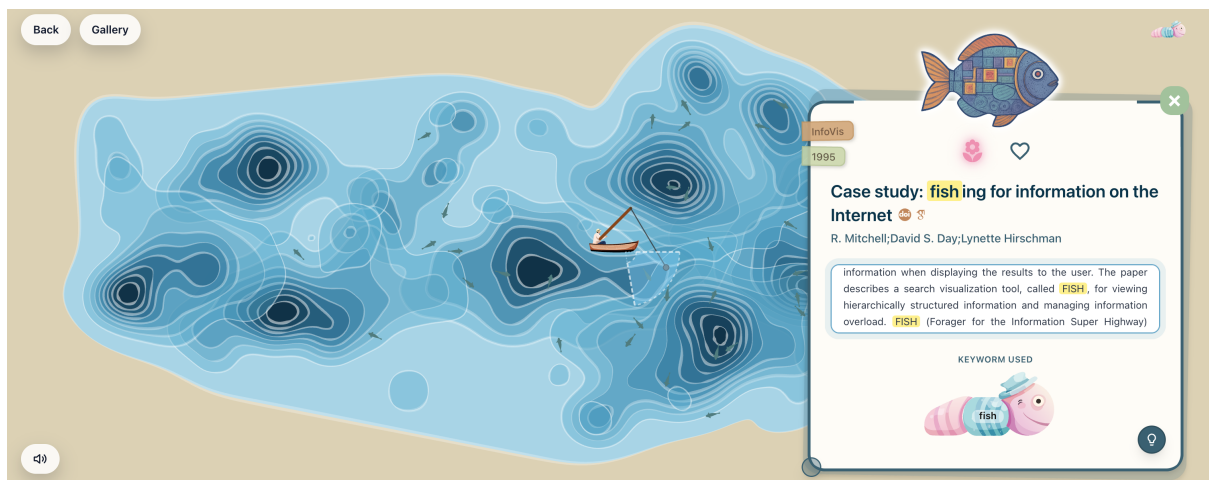


# Fishing for Papers: A Serendipitous Knowledge Discovery Game

J. Eschner<sup>1</sup>  and Y. Guo<sup>2</sup> 

<sup>1</sup>TU Wien, Austria

<sup>2</sup>Peking University, China



**Figure 1:** The Pond scene of *Fishing for Papers*, representing the paper pool of a selected venue. A collection of keywords, the Keyworm, is used as bait to catch papers. Papers matching a Keyworm are visualized as ripples on the pond with fish swimming around. The player steers a boat and casts a line to catch a paper. When a paper is caught, it can be saved to the gallery, represented by a generated image.

## Abstract

Fishing has long been used as a metaphor for information-seeking. In the context of literature research, we may see a query act as bait, with the researcher observing what papers “bite”. Taking this analogy literally, we created an open-ended fishing game for literature exploration. Our game aims to support serendipitous knowledge discovery in a playful, relaxing way. In our game, the player uses a virtual bait, the Keyworm (a set of keywords), to catch papers from a Pond (the paper pool for a given venue). To engage players in serendipitous discovery, we introduce a serendipity measurement that categorizes papers into varying levels of rarity. Players are encouraged to iteratively experiment with different keywords to navigate the paper corpus and catch unexpectedly interesting papers. Our game employs the cozy game design principles to provide a relaxing and enjoyable way for visualization researchers to explore the visualization literature. We evaluated the game in playtests with six researchers and made it available as an online demo: <https://fishing-for-papers.github.io/game/>.

## CCS Concepts

• **Human-centered computing** → **Visualization**; • **Applied computing** → **Computer games**; • **Information systems** → **Information retrieval**;

## 1. Introduction

The majority of information-seeking and knowledge-discovery systems aim to provide users with relevant, sought-after information in a structured, direct manner, as exemplified by Shneiderman’s Information-Seeking Mantra [Shn03]. However, this focused serv-

ing of information may lead to overspecialization and the creation of filter bubbles. In the context of recommender systems, this problem is known as the serendipity problem [DGLSM15]: in a highly focused information-seeking system, the user may miss surprisingly interesting items.

© 2026 The Author(s).

Proceedings published by Eurographics - The European Association for Computer Graphics. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

DOI: 10.2312/visgames.20261002

A study on serendipity in the information-seeking behavior of interdisciplinary scholars [FF03] highlighted that the phenomenon is an important component of knowledge discovery. We therefore aim to develop a game that supports visualization experts and researchers in the field in discovering unknown or unexpected visualization literature. However, the same principles could potentially be applied to explore other text corpora without major changes to the system.

By framing the information-seeking process as a game, we aim to shift the player’s focus away from the stress of finding useful results and towards enjoying the process itself. To achieve this shift, we implement our knowledge discovery game according to the principles of *cozy games*, a genre that emphasizes a safe, relaxing experience over competitive, goal-driven gaming. More formally, Krzywinska et al. [KBM\*25] define coziness in games along the dimensions of safety, abundance, and softness. *Safety* is signified by the absence of danger, there being no impending loss or threat, and activities being voluntary and without risk. *Abundance* indicates that nothing is lacking, pressing, or imminent. *Softness*, which is characterized by aesthetic signals that tell the player they are in a low-stress environment.

A common activity found in cozy games is fishing [KBM\*25, Xia23], as it is connotated with mindfulness and relaxation. Orthogonally to that, fishing also has a long history as a metaphor for search and information retrieval [Sut98, MDH95, LCG05, Wil90]. Hartel and Savolainen [HS16] designate fishing as one of 14 common pictorial metaphors for information. Therefore, a fishing game presents itself as a fitting setting for serendipitous knowledge discovery. The fishing metaphor also works for potential problems of serendipitous discovery [CP00]: you may catch more fish than you need, or at times, you may just catch an old, useless boot from the bottom of the lake.

By combining the principles of cozy games with the fishing metaphor for knowledge discovery, we develop a prototype of a game for visualization literature research. Restricting the domain to visualization, the game primarily targets experts and students seeking to explore this corpus. Our work makes the following contributions:

- A novel interactive interface for exploring visualization literature in a cozy setting
- A simple serendipity measurement for papers as a basis for gamification
- A publicly available demo game for serendipitous literature discovery: <https://fishing-for-papers.github.io/game/>.

## 2. Design Considerations

Inspired by the “fishing” metaphor of information foraging, our game integrates three key components: (1) a fishing mechanic that provides a playful, relaxing, and less overwhelming way to explore the literature; (2) visualization of the paper corpus and retrieval results that offers an aesthetic and structured view of the information space; (3) a serendipity-based reward system that guides players through iterative keyword exploration to discover rare and unexpected papers. In this section, we will detail the design considerations.

### 2.1. Goal and Principles

The core objective of our game is to transform the search for visualization literature into a playful, relaxing experience that encourages serendipitous discovery of interesting, sometimes niche, papers. As the goal is relaxation and discovery, it is designed to be open-ended and non-competitive. Apart from the constraints set through the interface, there are no additional rules to the game. Players are free to explore at their own pace, revisit previous exploration sessions, and search for new papers. While there is no traditional “winner”, the system employs a serendipity measure for each paper that players can use as an incentive to discover unusual or rare papers. Following the principles of safety and abundance in cozy games, there is no threat present in the game. The player cannot be harmed or lose the game, and is not under time pressure.

### 2.2. Aesthetics

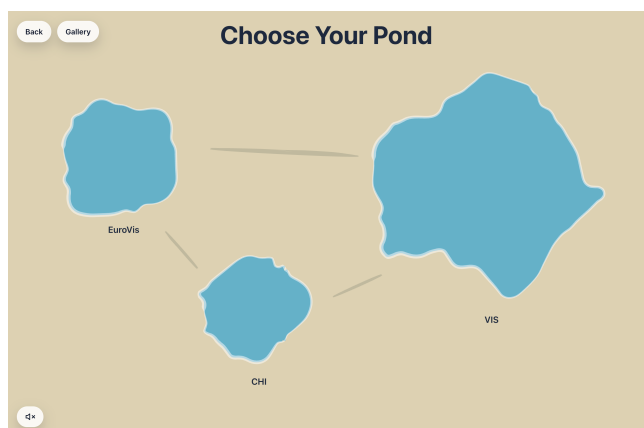
To craft a calm, relaxing environment in line with the softness principle in cozy games, we opted for a muted pastel color scheme paired with soothing background music. All visual assets, including the boat sprites and environment, are minimalist vector graphics with a hand-drawn style. This style is also reflected in the game’s menu screen, which features a vectorized image of a fishing boat on a serene mountain lake. The game’s soundscape is defined by a calm, looping background track titled “Search the World”, sourced from the OpenGameArt repository [Len15].

### 2.3. Visual Encoding

We employ a fishing metaphor throughout the game. Each individual paper is considered a fish, and the paper corpus forms a pond. Different publication venues are conceptualized as distinct ponds to explore. This mapping provides a visual structure for the papers and venues, transforming literature exploration into spatial navigation within and across ponds.

**Map of Venues.** As shown in Figure 2, the map scene presents different venues (IEEE VIS, EuroVis, and CHI) as a network of ponds arranged using a force-directed layout. We compute the pond layout using the d3 [BOH11] spring-embedded force simulation with 300 ticks. The node size of the pond encodes the paper count of each venue, while the shape corresponds to the outline of the semantic paper embeddings computed for each venue. The network is represented as a fully connected graph, where edges indicate paths the player takes from one pond to the next.

**Pond of Papers.** Each pond represents the corpus of papers from a venue, serving as the primary space where players explore and “fish” for papers. As shown in Figure 1, the pond’s visual design is inspired by natural water patterns, such as ripples, vortices, and depth variations. To provide an overview of paper distribution in the pond, we use contours to encode paper density based on their 2D embeddings, mimicking the water-depth encoding in bathymetric charts [EM69], i.e., darker regions indicate deeper areas with more papers. To achieve the aesthetic of bathymetric contour maps, we first cluster the papers into multiple groups and then compute




**Figure 2:** The Map of Ponds. Each pond represents the paper corpus of a venue (VIS, EuroVis, and CHI). The map is modeled as a network of ponds and is visualized using a force-directed layout.

density contours within each cluster. We computed semantic embeddings for each paper (combining title and abstract) using OpenAI’s `text-embedding-3-small` model [Ope26b]. The 2D embeddings of the papers are obtained by applying UMAP to the high-dimensional text embeddings; the clusters are computed using k-means on the high-dimensional embeddings.

**Ripples of Search Results.** In our game, players use a “bait” (a set of keywords) to fish for relevant papers via keyword-based retrieval. Inspired by the ripples formed on the water surface when fish move, we use ripple patterns to encode the spatial distribution of papers that match the specified keywords. For visual clarity, papers matching the keyword(s) are first grouped using a grid-based clustering method based on their 2D positions. For each group, the center of the ripple is determined by the average position of the papers within the cluster. Larger clusters with more papers are represented by larger ripples with more concentric rings. To further reinforce the pond metaphor and to visually emphasize the paper locations, we animate fish swimming around the center of each cluster. The number and size of the fish are proportional to the number of papers in each cluster.

**Gallery.** The final main design element of the game is the gallery view, where players can view and revisit their caught papers. Following the cozy games principles, the “fish” in the game are not caught and killed. Rather, whenever the player decides to keep a paper, a representative image of them is added to the gallery. To this end, the gallery view (see Figure 3) displays each captured paper as a picture frame with an image of the paper as a stylized fish. In order to get a unique image of each paper, we used text-to-image generative models, specifically OpenAI’s `gpt-image-1-mini` model [Ope26a], to generate an image for each saved paper. The prompt for image generation is constructed from the paper’s title and abstract, along with style modifiers such as “dreamy” and “flat art”, which align with the general aesthetics of our interface design.



**Figure 3:** The gallery of caught papers. Each frame represents a paper, with the icon beneath indicating its rarity level.

### 3. Fishing for Papers: Game

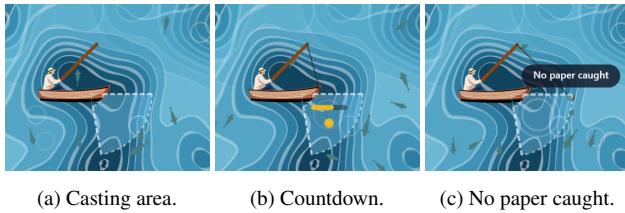
*Fishing for Papers* is an open-ended single-player game, designed to engage the player to seek and discover literature in a playful and relaxing way. The central element of the game is a fishing mechanic in which the player uses a virtual fishing rod to catch papers from the paper pool of a given venue. This fishing mechanic is inspired by single-player games such as *Animal Crossing: New Horizons* [Nin20] and dedicated fishing games like *Cast n Chill* [Bra25], which emphasize the relaxing process of fishing over competitive or score-based games.

The game is implemented as a browser-based PC game that requires only a desktop or laptop computer with internet access. The paper dataset (see Section 3.2) is hosted on the game server, requiring no additional preparation on the player’s part.

A typical session of the game consists of the following general steps: (i) starting a new literature research session, (ii) choosing a venue from the map of ponds to go fishing in, (iii) crafting a suitable bait “Keyword” for the desired topic, (iv) fishing papers from the pond, and (v) adding them to the gallery of papers. Once the player is happy with the results of their fishing, they can review the catches in the gallery and optionally download the collection of caught papers in BibTeX format.

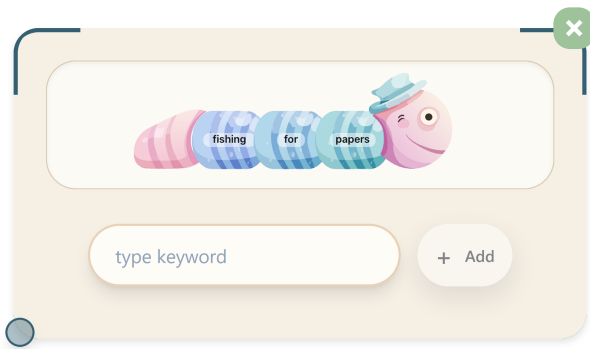
#### 3.1. Game Manual

The game starts on a menu screen where the player can resume an existing game (if available) or start a new one. When starting a new game, the player is presented with a map of different visualization venues as shown in Figure 2. Each venue, represented as a pond, is clickable and, when clicked, transports the player to a detailed pond view where the actual fishing game takes place. Here, the player steers a boat over the pond of paper embeddings using their WASD or arrow keys. When no input is provided via these keys, the boat will come to rest and switch from a top view to a side view, indicating it is now ready for casting a fishing line. However, just casting a line is not sufficient to catch a paper. As with actual fishing, some bait is required. The player crafts this bait by defining



**Figure 4:** Casting a line to catch a paper. The player clicks within the casting area (a) to cast the fishing line. Once the line hits the water, a timer starts to count down (b). If a paper was caught, an overlay window opens (see Figure 1); otherwise, a message informs the user that no paper was caught (c).

a so-called “Keyworm” from one or more keywords. This is done by clicking the Keyworm icon in the top-right corner, which opens an editor for adding and removing keywords (Figure 5).



**Figure 5:** The keyword interface allows the player to craft their keyword query, which serves as bait for the fishing.

As soon as at least one keyword is defined and matching literature is present in the dataset, some fish swarms and bubbles will appear on the map, indicating areas that may be promising for catching a relevant paper. The player now sails their boat to such an area and casts a fishing line. Casting is performed by clicking anywhere within the cast area next to the boat (see Figure 4). To successfully catch a paper, the fishing line should be positioned as close as possible to the center of the fish swarm or bubble that the player is aiming for. Once the line has been cast, a progress bar will appear, counting down until a potential bite. If there is a paper in close vicinity to the cast position, it will now be shown to the player in an overlay window. If none were found in the area, a message appears stating that no paper was caught.

For every paper caught by the player, they can review its title and abstract and navigate to the full version via its DOI. If deemed relevant or interesting, they add it to their paper gallery by clicking the heart icon. An additional icon indicates how rare and unexpected each caught paper is (see Section 3.4). To gain insights into how the serendipity of a paper is calculated, a light bulb icon at the bottom of the paper record lets the player highlight the keywords that contributed to the serendipity score.

If the player decides not to keep the paper, they release it back

into the pond by closing the overlay window. Every paper added to the gallery receives a unique, generated fish image, which will also serve as its visual representation.

In a single fishing session, the player can define an arbitrary number of Keyworms and explore all available venues, collecting as many papers as they want. Instead of obtaining an overall score for the caught papers, the only “success metric” is the different levels of papers, which are also displayed in the gallery.

### 3.2. Dataset

The design of our game works with any paper dataset. Here, we choose to start with the VisPubs dataset [Lan24], as it is well-curated and most relevant to the visualization community. The VisPubs dataset is a collection of visualization and HCI publications spanning 1986 to 2025, comprising 4,008 VIS papers (including InfoVis, SciVis, and VAST), 1,197 EuroVis papers, and a subset of 887 visualization-related CHI papers.

### 3.3. Keyword Matching

As described in Section 3.1, the player uses a Keyworm, which consists of a set of keywords, to retrieve relevant papers. The retrieval is performed using a conjunctive (AND) logic, where all keywords must be present, and is implemented via *exact* substring matching without boundary checking. To reduce visual clutter, if the number of matched papers exceeds a threshold ( $T = 100$ ), the system randomly sample  $T$  papers from the retrieval results to present to the user. Since exact matching can be restrictive, i.e., specific queries with multiple keywords may return no results, we employ BM25 retrieval as a fallback strategy. When no exact match is found, the system presents the top- $k$  BM25 results ( $k = 10$ ).

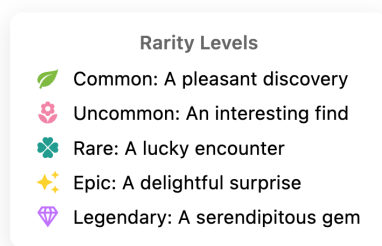
### 3.4. Reward System

Our game is designed to encourage serendipitous discovery of papers in a relaxing way. According to Watson, serendipity can be defined as “a happy accident in which an information seeker unexpectedly stumbled across relevant information” [Wat08]. In our game, rather than optimizing for efficiency or targeted search, we invite players to wander, experiment with different keyword combinations, and uncover interesting papers that might be unusual, rare, or difficult to find through conventional search approaches. To achieve this, we measured each paper’s serendipity. Players are encouraged to catch papers that have high levels of serendipity.

According to Kotkov et al. [KWV16], serendipity can be characterized by relevance, novelty, and unexpectedness. Among these, we focus on the unexpectedness dimension, which is quantified using a mutual information-based approach, while the other dimensions are more subjective, and we leave them to user perception. To operationalize unexpectedness, we measure the rarity of keyword combinations in each paper. We first extract the top-10 keywords for each paper using TF-IDF (Term Frequency-Inverse Document Frequency). We then generate all possible 2-way combinations from these keywords and calculate their rarity using negative pointwise mutual information (negative PMI) [CH90]. Unlike simple frequency-based measures, negative PMI quantifies how much

less two keywords co-occur than would be expected by chance:  $\text{neg-PMI} = \log \frac{P(A) \cdot P(B)}{P(A,B)}$ , where  $P(A)$  and  $P(B)$  are the probabilities of individual keywords appearing, and  $P(A,B)$  is the probability of their co-occurrence. This approach ensures that high scores are assigned to truly unexpected combinations, i.e., high-frequency keywords that rarely appear together, rather than simply rare keywords. For example, two common topics that surprisingly avoid each other in the literature would score higher than two obscure topics that rarely co-occur simply because both are uncommon.

We aggregate these combination scores to produce a raw unexpectedness score for each paper, which is then mapped to five rarity levels (Common, Uncommon, Rare, Epic, Legendary) using quantile-based mapping. The mapping is designed to create a left-skewed distribution where 40% of papers are classified as Common, 30% as Uncommon, 15% as Rare, 10% as Epic, and only 5% achieve the Legendary status, making high-rarity discoveries feel genuinely special and serendipitous in the game experience (Figure 6). To improve transparency of the computation process and provide players with hints on how to find rare papers, players can choose to highlight the keywords involved in the metric calculation by toggling the hint button at the bottom right of the paper information box (Figure 1).



**Figure 6:** Five levels of rarity derived from the serendipity measurement. The rarity level here serves only as a proxy for serendipity. It is designed to encourage exploration and does not reflect any assessment of paper quality or values.

## 4. Reflection & Conclusion

In this section, we report on our pilot playtests, in which six researchers conducted short, supervised literature research sessions using our game to give feedback and suggestions for further improvement. We then reflect on those findings before concluding with an outlook on potential future work.

### 4.1. Playtests

To evaluate the game's suitability for serendipitous knowledge discovery, we conducted informal, qualitative playtests with six researchers, recruited from within our labs. Four of our test players were visualization researchers, one was a computer vision researcher with an interest in visualization, and one was an educational game developer. Playtest participants provided verbal consent to publish anonymized and summarized reports of their sessions, and, given the evaluation's informal nature, no ethics approval was required. After an initial introduction to the interface

and controls, the players were left to freely explore the game under the authors' observation. They were encouraged to perform *thinking aloud*, and the authors took notes on how the game was played, which keywords were used, and what problems occurred. Based on these records, we conducted an open coding round to derive generalizable findings. Note that the limited number of participants and the qualitative setting of the playtests merely constitute a feedback round, and a future user study would be required to make definitive statements on the game's effectiveness.

In the playtests, we observed a range of exploratory behaviors that highlight both the system's potential for serendipitous discovery and areas for improvement. Players frequently began with either targeted keyword searches related to their own research topics or more open-ended, playful strategies, often refining keywords iteratively based on previous catches. In one case, a player combined "geo" and "reasoning" to locate a relevant paper, inferred "mobile" as a distinguishing keyword, and immediately obtained a rare ("legendary") result. In another instance, a player encountered a paper closely related to a colleague's work and shared it, demonstrating the game's capacity to facilitate unexpected yet meaningful connections. As players were not instructed to perform a specific type of literature research, we observed different patterns, where some players stayed with a main keyword (e.g., "treemap") and then explored combinations with additional keywords, while others varied their keywords to explore as much of the pond as possible.

Some players also adopted self-defined goals beyond literature search, such as generating increasingly unusual keyword combinations or exploring the visual diversity of generated fish images. One of those players reported becoming deeply immersed, even losing track of time during the session.

At the same time, some usability challenges were identified, particularly regarding precisely navigating the boat to the desired spot (three players) and targeting the fishing line casts (two players). Players also suggested potential future enhancements, such as incentive systems (e.g., looking for specific papers or paper types) or the inclusion of social interactions (e.g., showing other players on the pond). Overall, these initial playtests indicate that the game supports serendipitous knowledge discovery while engaging players to keep exploring.

### 4.2. Summary

Across sessions, we observed that enjoyment emerged not from a single objective, but from a combination of complementary motivations. We identified three main factors in the game that contribute to making a fun and engaging session. First, aiming at *catching a rare ("legendary") paper* encouraged players to try different keywords to explore the corpus. This rarity-based incentive structure provided a clear, game-like goal that sustained curiosity and experimentation. Second, players reported satisfaction when *encountering papers that were relevant to their own work*, even if these were not of the highest rarity. Third, by creating a unique image for each saved paper, our game incentivizes players to add papers to their library and to discover visually interesting fish (i.e., testing what a fish for a certain type of paper would look like).

Based on our observations, we consider a play session success-

ful when it enables at least one of the following outcomes: (1) the discovery of unexpected or unusual papers, (2) the compilation of a useful set of literature for further research, or (3) a state of sustained engagement or immersion during exploration. Importantly, these criteria highlight that success is not solely tied to efficiency, but also to the serendipitous nature of the exploratory experience.

### 4.3. Future Work

The findings from our initial evaluation open several directions for future development. First, improving onboarding and guidance represents a key step toward accessibility. Tutorials, visual cues, or adaptive hints could help players better understand core mechanics such as navigation, casting, and keyword construction.

The underlying search mechanism could be extended beyond keyword-based queries and simple string matching. Incorporating semantic or vector-based search techniques may enable more robust and flexible exploration, allowing players to discover relevant papers even when keywords are too broad or ambiguous.

It is easy to sustain the current game, i.e., adding new visualization publications, by updating the embeddings. The existing projection model can be reused to transform new papers into the current projection, and a new map can be built every few years when there are drastic shifts in the research landscape. Additionally, the system could be expanded to integrate larger, more diverse corpora. Automating data acquisition through APIs and generalizing beyond manually curated venues would increase coverage and open the game to other scientific areas. To support sustainability and community contributions, we release the game as open source at: <https://github.com/fishing-for-papers/game>.

Finally, several design extensions could further enhance engagement. These include progression systems (e.g., achievements or curated collections), social features (e.g., shared discoveries or visible co-exploration), support for annotating explored knowledge, and alternative platforms (e.g., mobile devices) to support more casual, on-the-go interaction.

### Acknowledgments

This research was funded in whole or in part by the Austrian Science Fund (FWF) [10.55776/P36453](https://www.fwf.ac.at/en/foyer/10.55776/P36453). This work was supported by the China Scholarship Council (CSC).

### References

- [BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D<sup>3</sup> Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (Dec. 2011), 2301–2309. URL: <https://ieeexplore.ieee.org/abstract/document/6064996>, doi:10.1109/TVCG.2011.185. 2
- [Bra25] BRAWLER W.: Cast n Chill, 2025. URL: <https://www.wombatbrawler.com/cast-n-chill>. 3
- [CH90] CHURCH K., HANKS P.: Word association norms, mutual information, and lexicography. *Computational linguistics* 16, 1 (1990), 22–29. 4
- [CP00] COOPER J., PRAGER J.: Anti-serendipity: finding useless documents and similar documents. In *Proceedings of the 33rd Annual Hawaii International Conference on System Sciences* (Jan. 2000), pp. 8 pp. vol.1–. URL: <https://ieeexplore.ieee.org/document/926691>, doi:10.1109/HICSS.2000.926691. 2
- [DGLSM15] DE GEMMIS M., LOPS P., SEMERARO G., MUSTO C.: An investigation on the serendipity problem in recommender systems. *Information Processing & Management* 51, 5 (Sept. 2015), 695–717. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0306457315000837>, doi:10.1016/j.ipm.2015.06.008. 1
- [EM69] EDVALSON F. M., MISCOSKI V. T.: *Bathymetric Charts: Their Development and Use*. U. S. Naval Oceanographic Office, 1969. 2
- [FF03] FOSTER A., FORD N.: Serendipity and information seeking: an empirical study. *Journal of Documentation* 59, 3 (June 2003), 321–340. URL: <https://doi.org/10.1108/00220410310472518>, doi:10.1108/00220410310472518. 2
- [HS16] HARTEL J., SAVOLAINEN R.: Pictorial metaphors for information. *Journal of Documentation* 72, 5 (Sept. 2016), 794–812. URL: <https://doi.org/10.1108/JD-07-2015-0080>, doi:10.1108/JD-07-2015-0080. 2
- [KBM\*25] KRZYWINSKA T., BROWN D., MA M., BELINSKIY A., BHUI K.: Coziness in Games: Second Homes, Audiences, and Aesthetics. *Games and Culture* (Jan. 2025), 15554120241310920. URL: <https://doi.org/10.1177/15554120241310920>, doi:10.1177/15554120241310920. 2
- [KWV16] KOTKOV D., WANG S., VEIJALAINEN J.: A survey of serendipity in recommender systems. *Knowledge-Based Systems* 111 (2016), 180–192. 4
- [Lan24] LANGE D.: Vispubs.com: A visualization publications repository, 2024. doi:10.31219/osf.io/dg3p2. 4
- [LCG05] LUO F.-F., CHEN G.-L., GUO W.-Z.: An improved "fish-search" algorithm for information retrieval. In *2005 International Conference on Natural Language Processing and Knowledge Engineering* (Oct. 2005), pp. 523–528. URL: <https://ieeexplore.ieee.org/abstract/document/1598793>, doi:10.1109/NLPKE.2005.1598793. 2
- [Len15] LENTZ T.: Search the world. <https://opengameart.org/content/search-the-world-trevor-lentz>, 2015. Audio file, licensed under CC-BY 3.0. 2
- [MDH95] MITCHELL R., DAY D., HIRSCHMAN L.: Case study: fishing for information on the Internet. In *Proceedings of Visualization 1995 Conference* (Oct. 1995), pp. 105–111. URL: <https://ieeexplore.ieee.org/abstract/document/528693>, doi:10.1109/INFVIS.1995.528693. 2
- [Nin20] NINTENDO: Animal Crossing: New Horizons – Official site, 2020. URL: <https://animalcrossing.nintendo.com/new-horizons/>. 3
- [Ope26a] OPENAI: gpt-image-1-mini. <https://developers.openai.com/api/docs/models/gpt-image-1-mini>, 2026. Accessed: March 2026. 3
- [Ope26b] OPENAI: text-embedding-3-small. <https://developers.openai.com/api/docs/models/text-embedding-3-small>, 2026. Accessed: March 2026. 3
- [Shn03] SHNEIDERMAN B.: The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *The Craft of Information Visualization*. Elsevier, 2003, pp. 364–371. URL: <https://linkinghub.elsevier.com/retrieve/pii/B9781558609150500469>, doi:10.1016/B978-155860915-0/50046-9. 1
- [Sut98] SUTCLIFFE A.: Task-related navigation in information spaces. *Workshop on personalized and social navigation in information space* (1998). 2
- [Wat08] WATSON E. A.: Going fishing: Serendipity in library and information science. *University of North Carolina at Chapel Hill* (2008). 4
- [Wil90] WILSON J. A.: Fishing for Knowledge. *Land Economics* 66, 1 (Feb. 1990), 12. URL: <http://www.jstor.org/stable/3146679?origin=crossref>, doi:10.2307/3146679. 2
- [Xia23] XIAO L. Y.: What's a Mini-Game? The Anatomy of Fishing Mini-Games. In *Conference Proceedings of DiGRA 2023 Conference: Limits and Margins of Games Settings* (June 2023). URL: <https://dl.digra.org/index.php/dl/article/view/1911>, doi:10.26503/dl.v2023i1.1911. 2