

An Affordance-Based Framework for Human Computation and Human-Computer Collaboration

Category: Research

ABSTRACT

Visual Analytics is “the science of analytical reasoning facilitated by visual interactive interfaces” [70]. The goal of this field is to develop tools and methodologies for approaching problems whose size and complexity render them intractable without the close coupling of both human and machine analysis. Researchers have explored this coupling in many venues: VAST, Vis, InfoVis, CHI, KDD, IUI, and more. While there have been myriad promising examples of human-computer collaboration, there exists no common language for comparing systems or describing the benefits afforded by designing for such collaboration. We argue that this area would benefit significantly from consensus about the design attributes that define and distinguish existing techniques. In this work, we have reviewed 1,271 papers from many of the top-ranking conferences in visual analytics, human-computer interaction, and visualization. From these, we have identified 49 papers that are representative of the study of human-computer collaborative problem-solving, and provide a thorough overview of the current state-of-the-art. Our analysis has uncovered key patterns of design hinging on human- and machine-intelligence affordances, and also indicates unexplored avenues in the study of this area. The results of this analysis provide a common framework for understanding these seemingly disparate branches of inquiry, which we hope will motivate future work in the field.

Index Terms: Human computation, human complexity, theory, framework.

1 INTRODUCTION

Thomas and Cook define the field of Visual Analytics as “the science of analytical reasoning facilitated by visual interactive interfaces” [70]. By leveraging increasing computational power and the significant bandwidth of human visual processing channels, it strives to facilitate the analytical reasoning process and support the “human capacity to perceive, understand, and reason about complex and dynamic data and situations” [70]. As the field matures, it is increasingly imperative to provide mechanisms for approaching analytic tasks whose size and complexity render them intractable without the close coupling and dynamic interplay of both human and machine analysis. Primary goals of this field are to develop tools and methodologies that facilitate human-machine collaborative problem solving, and to understand and maximize the benefits of such a partnership.

Researchers have explored this coupling in many venues: IEEE Conference on Visual Analytics Science and Technology (VAST), IEEE Visualization Conference (Vis), IEEE Information Visualization Conference (InfoVis), ACM Conference on Human Factors in Computing Systems (CHI), ACM Conference on Knowledge Discovery and Data Mining (KDD), ACM Conference in Intelligent User Interfaces (IUI), and more. The study of general human-computer collaboration offers a plethora of examples of successful human/machine teams [15, 23, 37, 39, 41, 46, 50, 65, 66, 68, 83]. Developments in supervised machine learning in the visualization community present several vetted techniques for human intervention into computationally complex tasks [3, 4, 12, 18, 28, 29, 36, 47, 51, 57]. The emerging field of human computation inverts the traditional paradigm of machines providing computational support

for problems that humans find challenging, and demonstrates success using aggregated human processing power facilitated by machines to perform difficult computational tasks such as image labeling [21, 33, 73, 74], annotating audio clips [44, 49], and even folding proteins [20]. While there have been a multitude of promising examples of human-computer collaboration, there exists no common language for describing such partnerships. This begs several questions:

1.1 Problem selection

How do we tell if a problem would benefit from a collaborative technique? Balancing the cost of building and deploying a collaborative system with the benefits afforded by its use is currently precarious at best. Without a framework in which to situate the development of new systems, we rely heavily on researcher intuition and current fieldwide trends to decide which problems to approach using these techniques. This is akin to looking for the sharpest needle in a haystack of needles, and while it has led to many novel approaches to hard problems, it has also led to the investment of significant time and energy into inefficient collaborative solutions for problems that might better have been (or have already been) solved by human or machine techniques alone.

1.2 Function allocation

How do we decide which tasks to delegate to which party, and when? It has long been stated (even by the author himself) that Fitts’ HABA-MABA lists [27] are insufficient and out-of-date. Sheridan notes that function allocation in collaborative systems is far from a perfect science [63]. Dekker argues that static function allocation consistently misses the mark because humans adapt to their surroundings, including systems with which they work [22]. However, the effectiveness of any collaborative system is deeply rooted in its ability to leverage the best that both humans and machine have to offer. Because of this, a language for describing the skills and capacity of the collaborating team is long overdue.

1.3 Comparative analysis

Finally, how does one system compare to others trying to solve the same problem? With no common language or measures by which to describe new systems, we must rely heavily on researcher intuition and anecdotal evidence. This makes it challenging to reproduce results and to build on previous discoveries, leading to the development of many one-off solutions rather than a cohesive, directed line of research.

We argue that each of these areas would benefit from consensus about the set of attributes that define and distinguish existing techniques. In this work, we have reviewed 1,271 papers from many of the top-ranking conferences in visual analytics, human-computer interaction, and related areas. From this corpus, we have identified 49 that we believe are representative of the study of human-computer collaborative problem-solving, and provide a thorough overview of the current state-of-the-art. Our analysis has uncovered key patterns of design hinging on human- and machine-intelligence affordances, or properties of the human and machine collaborators that offer opportunities for collaborative action. The results of this analysis provide a common framework

for understanding these seemingly disparate branches of research and also indicate unexplored avenues in the study of this area, which we hope will motivate future work in the field.

1.4 Contributions

This paper makes several contributions to the study of human-computer collaboration and human computation in Visual Analytics and HCI:

- First, we conduct a comprehensive literature review of 1,271 papers in this area and present a collection of 49 publications that represent the state of the art in the study of human-computer collaboration and human computation.
- Second, from these papers we identify groupings based on human- and machine-intelligence affordances. These groupings form the basis of a common framework for understanding and discussing this collection of work.
- Third, we identify unexplored areas for future work and argue for the utility of this framework in addressing the questions outlined in the previous section.

2 BACKGROUND

Before we present our survey and associated framework, let us first briefly describe two terms that will be utilized extensively throughout this paper: *human-computer collaboration* and *human computation*.

2.1 Human-Computer Collaboration

In a 1993 symposium at AAAI, researchers from a variety of backgrounds came together to discuss challenges and benefits in the emerging field of human-computer collaboration. They defined collaboration as *a process in which two or more agents work together to achieve shared goals*, and human-computer collaboration as *collaboration involving at least one human and at least one computational agent* [69]. This collaboration has also been called *mixed-initiative systems* [34], in which either the system or the user can initiate action, access information and suggest or enact responses [70]. Mixed-initiative systems have been explored in diverse areas including knowledge discovery [71], problem-solving in AI [26], procedural training in virtual reality [56] and much more. The field of Visual Analytics is deeply rooted in human-computer collaboration; that is, Visual Analytics seeks to leverage both analyst intelligence and machine computational ability in a collaborative effort to solve complex problems.

2.2 Human Computation

In his 2005 doctoral thesis [72], Luis von Ahn introduces the concept of *human computation*; that is, harnessing human time and energy for solving problems that have to date proven computationally intractable. This is accomplished by treating human brains as processors in a distributed system. With the advent of online marketplaces providing an on-demand workforce for microtasks, we have seen an explosion of work utilizing human processing power to approach problems such as image labeling [21, 33, 73, 74], digitizing text [77], annotating audio clips [44, 49], and folding proteins [20]. For surveys of research in this area, please see [55, 82].

While much of the work in this has shown promise in harnessing human computational power, there is a temptation to use human workers as an easy out. In his article entitled “Why I Hate Mechanical Turk Research (and Workshops)” [1], Adar argues:

We should not fool ourselves into believing that all hard problems fit this mold [warrant human computation] or completely distract ourselves from advancing other, computational means of solving these problems. More

importantly, we should not fool ourselves into believing that we have done something new by using human labor. . . Showing that humans can do human work is not a contribution.

This sentiment has prompted fascinating debate about when and how to leverage human intelligence in computation.

2.3 Previous Frameworks

A few of the existing papers surveying work in Human-Computer Collaboration and Human Computation also include discussions of the design dimensions that organize and contextualize their work. In these surveys, the authors provide mechanisms to compare and contrast the systems they review to others along salient dimensions.

Bertini and Lalanne [9] survey the intersection of machine learning and visualization, identifying three categories of design hinging on the distribution of labor between human and machine. In *enhanced visualization*, human use of the visualization is the primary data analysis mechanism and automatic computation provides additional support in the form of *projection*, *intelligent data reduction*, and *pattern disclosure*. In *enhanced mining*, data analysis is primarily accomplished by the machine through data mining and visualization provides an advanced interactive interface to help interpret the results through *model presentation* and *patterns exploration and filtering*. In *integrated visualization and mining*, work is distributed equally between the human and machine collaborators at different stages: *white-box integration*, where the human and machine cooperate during model-building, and *black-box integration*, where the human is permitted to modify parameters of the algorithm and immediately visualize the results.

In the area of human computation, Yuen et al. [82] identify three broad categories based on the relative maturity of the system. *Initiatory* systems are the earliest examples of human computation and were generally used to collect commonsense knowledge. *Distributed* systems were the next generation of human computation, aggregating the contributions of Internet users but with limited scalability and without any mechanism to guarantee the accuracy of the information collected. Finally, the authors describe *social game-based* systems, the most recent incarnation of human computation involving enjoyable, scalable and reliable systems for approaching hard AI problems. In a later survey, Quinn and Bederson [55] identify six dimensions along which they characterize human computation systems. *Motivation* describes the mechanism for encouraging human participation. *Quality control* indicates whether and how a quality standard is enforced upon the human workers. *Aggregation* refers to the means by which human contributions are collected and used to solve the problem at hand. The remaining dimensions of *human skill*, *process order*, and *task-request cardinality* are self-explanatory.

Each of these frameworks provides critical insight into organizing the systems appearing in the venues they survey. However, because each is specific to a specific subclass of collaborative systems, it is difficult to extend them to a broader class of human-computer collaborative systems. In the following sections, we provide a detailed survey of the literature across many venues, and argue for examining these systems through the lens of *affordances*; that is, what does each collaborator (machine or human) bring to the table in support of the shared goals of the team?

3 FRAMEWORK: ALLOCATION AND AFFORDANCES

We now introduce the foundation upon which we will build our framework for describing and understanding human-computer collaborative systems.

3.1 Function allocation in human-machine systems

Researchers have sought a systematic approach for the appropriate allocation of functions to humans and machines for decades.

In 1950, Fitts published the first formal attempt to characterize functions performed better by machines than humans, and vice versa [27]. For years, this list was regarded as the definitive mantra for function allocation, despite the author’s assertion that to use his list to determine function allocation was to lose sight of the most basic tenet of a human-machine collaborative system. As later articulated by Jordan, this underlying foundation is that humans and machines are *complementary*, rather than antithetical [38]. Price [54] further expands on this idea by arguing that function allocation is perhaps better envisioned as an iterative process rather than a decisive listing, and that there may be more than one optimal allocation for any given problem. Price also notes that human operators require support to perform optimally, and emphasizes the importance of understanding cognitive loading and engagement.

In more recent work, several contemporaries have argued that the notion of function allocation as it was originally conceived no longer makes sense. Sheridan discusses several problems with function allocation which include ever-increasing computing power, complicated problems with optimal allocation differing at each stage, and ill-defined problem spaces [63]. Dekker and Woods provide a second counterargument to the validity of any Fitts-style HABA-MABA listing in [22]. They point out a relationship that is often leveraged (though seldom explicitly stated) by the field of Visual Analytics: human-machine collaboration transforms human practice and forces people to adapt their skills and analytic practices. They argue for a shift in attention, moving away from allocation of tasks to a focus centered on how to design for harmonious human-machine cooperation. That is, how do we get humans and machines to play nicely, and work effectively?

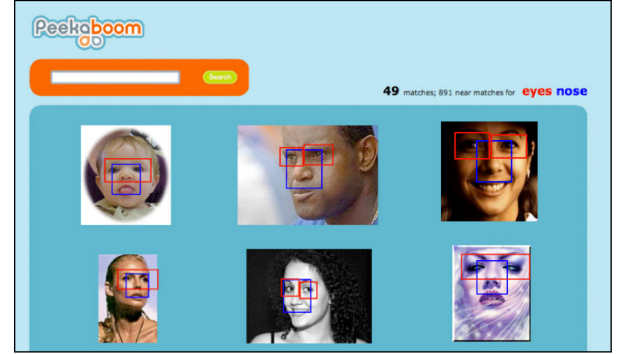
3.2 Affordances

In 1977, American psychologist J.J. Gibson stated his theory that an organism and its environment complement each other [31], which is much in alignment with the work by Jordan cited in the previous section. In this work, Gibson coined the term *affordances*, defining them as the opportunities for action provided to an organism by an object or environment. Norman later appropriated this term as it applies to design and the field of Human-Computer Interaction, redefining it slightly to refer only to the action possibilities that are readily perceivable by a human operator [53]. This definition shifts the concept of affordance toward *relational* rather than subjective or intrinsic; that is, an affordance exists *between* an actor and the object or environment, not existing separate from that relationship.

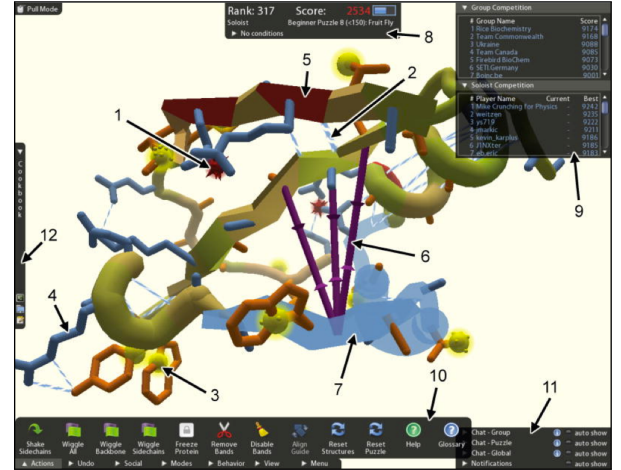
In the case of human-computer collaboration, we argue that there exist affordances in both directions. Both human and machine bring to the partnership opportunities for action, and each must be able to perceive and access these opportunities in order for them to be effectively leveraged. These affordances define the interaction possibilities of the team, and determine the degree to which each party’s skills can be utilized during collaborative problem-solving. In the next sections, we will survey the existing literature through the lens of affordances, providing a common framework for understanding and comparing research in the areas of human-computer collaboration, human intervention, and human computation. The affordances we identify are by no means an exhaustive list; they represent the patterns of design that we have seen in the existing literature of an emerging area. Please note that while examples will generally be given under the heading of a single affordance, systems mentioned may utilize multiple affordances (both human and machine) at the same time. For a complete listing of the affordances identified in all systems surveyed, please see Table 1 in the Appendix of this work. In Section 6, we present case studies of specific systems to discuss the costs and benefits of leveraging multiple affordances.

4 HUMAN AFFORDANCES

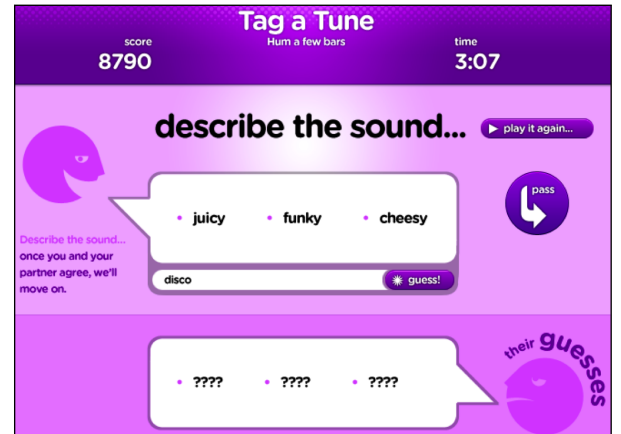
The human-computation and human-computer collaborative systems we have reviewed leverage a variety of skills and abilities afforded by the human participants. In this section, we will offer a brief definition of each of the affordances we have observed in the literature, discuss the utility of these affordances as articulated in the work reviewed and offer an overview of the application of each affordance in existing work.



(a) Peekaboom [76]



(b) Fold.it [20]



(c) TagATune [44]

Figure 1: Systems leveraging human affordances: (a) Visual perception, (b) Visuospatial thinking, and (c) Audiolinguistic ability.

4.1 Visual perception

Of the human affordances we will discuss, perhaps the most salient to the study of Visual Analytics is *visual perception*¹. In [64], Shneiderman comments on humans' capacity for visual processing:

[T]he bandwidth of information presentation is potentially higher in the visual domain than for media reaching any of the other senses. Humans have remarkable perceptual abilities... Users can scan, recognize, and recall images rapidly, and can detect changes in size, shape, color, movement, or texture. They can point to a single pixel, even in a megapixel display, and can drag one object to another to perform an action.

Given its direct applicability, it is perhaps unsurprising that we have seen a plethora of work in Visual Analytics and HCI leveraging human visual processing. For example, human visual perceptive abilities are utilized by Peekaboom [76] to augment image labels on the web (see Fig. 1a). For some tasks such as image labeling [21, 33, 59, 62, 73, 74], visual search [6, 10], and query validation [48, 80], the systems presented rely heavily on the users' visual perceptive abilities, with the machine serving only as a facilitator between the human and the data. For other tasks such as exploring high-dimensional datasets [68, 83], classification [4, 51], and dimension reduction [28, 36], machine affordances (which will be discussed at length in Section 5) are combined with human visual processing to achieve superior results.

4.2 Visuospatial thinking

A level deeper than basic visual processing such as image recognition, another skill afforded by human collaborators is *visuospatial thinking*², or our ability to visualize and reason about the spatial relationships of objects in an image. These abilities are strongly informed by our experiences in the physical world, which shape our understanding and are intrinsic to our everyday lives. We are able to visualize complex spatial relationships and tune this attention to accomplish specific goals. In an article on the significance of visuospatial representation in human cognition [60], Tversky notes:

For human cognition, [entities] are located in space with respect to a reference frame or reference objects that vary with the role of the space in thought or behavior. Which things, which references, which perspective depend on the function of those entities in context... These mental spaces do not seem to be simple internalizations of external spaces like images; rather, they are selective reconstructions, designed for certain ends.

We have seen evidence that progress can be made on computationally intractable problems through the application of human visuospatial thinking. For example, the Fold.it project (see Fig. 1b) has demonstrated remarkable success at protein folding [20], a problem known to be NP-complete [7] using purely computational means.

4.3 Audiolingistic ability

Another affordance presented by the human user is *audiolingistic ability*; that is, our ability to process sound³ and language⁴. Although separate from the visual affordances generally leveraged in Visual Analytics systems, we suggest that the interplay between visual and nonvisual human faculties is equally important in supporting analytical reasoning. In [70], Thomas and Cook state:

We perceive the repercussions of our actions, which also recalibrates perception, ensuring that vision, hearing, and touch maintain their agreement with each other. If we are to build richly interactive environments that aid cognitive processing, we must understand not only the levels of perception and cognition but also the framework that ties them together in a dynamic loop of enactive, or action-driven, cognition that is the cognitive architecture of human-information processing.

The literature contains several examples of systems leveraging this affordance. The well-known reCAPTCHA [77] system uses human linguistic ability augment computer vision in an effort to fully digitize the world's libraries. In MonoTrans2 [35], it is used to improve automated translation results using monolingual translators. In TagATune [44], human audio processing ability is leveraged to generate descriptive tags for music clips (see Fig. 1c). We have also surveyed examples utilizing human audio linguistic ability for audio annotation [5, 44, 49], transcription [16], and even crowd-sourced word processing [8].

4.4 Sociocultural awareness

In addition to physical senses, human collaborators also afford attributes such as *sociocultural awareness*, which refers to an individual's understanding of their actions in relation to others and to the social, cultural, and historical context in which they are carried out. Researchers in the area of embodied interaction have long advocated for design that acknowledges the importance of this relationship. In [24], Dourish notes:

[O]ur daily experience is social as well as physical. We interact daily with other people, and we live in a world that is socially constructed. Elements of our daily experience – *family, technology, highway, invention, child, store, politician* – gain their meaning from the network of social interactions in which they figure. So, the social and the physical are intertwined and inescapable aspects of our everyday experiences.

We argue that this can be viewed as an affordance, not just a complicating factor. For example, in Mars Escape [17], human participants partner with a virtual robot to complete collaborative tasks to build robust social training datasets for human-robot interaction research. This affordance is integral to the construction of common-sense knowledge databases [43, 45, 75], and has been leveraged in domains such as stress relief [19] and providing social scripts to support children with autism [11].

4.5 Creativity

Another important affordance of human collaborators is *creativity*⁵. As noted by Fitts [27], Dekker [22] and many others, humans are capable of incredible creativity, generating spontaneous arrhythmic approaches to problems that may be difficult or impossible to simulate. In [58], American psychologist Mark Runco posits:

It may be that creativity plays a role in all that is human. This surely sounds like a grand claim, but consider how frequently we use language or are faced with a problem. Think also how often problems are subtle and ill-defined... [C]reativity plays a role in each of our lives, and it does so very frequently.

We have seen human creativity leveraged to great success in both physical and conceptual design. For example, Yu and Nickerson [81] use human creativity to crowdsource design sketches via a human genetic algorithm, and Tanaka et al. [67] use sequential

¹For more on visual perception, see Gibson [32].

²For more on visuospatial thinking, see Shah and Miyake [60].

³For more on psychoacoustics, see Fastl and Zwicker [25].

⁴For more on language, see Vygotsky [78].

⁵For more on creativity, see Amabile [2].

application of crowds to produce creative solutions for social problems. Creativity has also been used to augment automated systems and find hidden outliers [47].

4.6 Domain knowledge

The final example of human affordance that we have seen in the literature is straightforward, but worthy of inclusion nonetheless. This is the affordance of *domain knowledge*. In their 2009 article on Knowledge-Assisted Visualization [13], Chen et al. argue:

[T]he knowledge of the user is an indispensable part of visualization. For instance, the user may assign specific colors to different objects in visualization according to certain domain knowledge. The user may choose certain viewing positions because the visualization results can reveal more meaningful information or a more problematic scenario that requires further investigation.

Often, this domain knowledge can be difficult or impossible to embed fully into the system itself, or it may be too time-consuming to generate a complete model of the domain. Instead, we can leverage the experience of the human analyst as part of the collaborative process. For example, we have seen domain expertise leveraged to help diagnose network faults [47], classify MRI data [12], perform domain-specific data transformations [40], and navigation and infer trends about a specific geographic region [4].

5 MACHINE AFFORDANCES

For over two decades, the HCI community has been engaged in conversation about affordances in technology, beginning with [30]. While much of the focus has centered on designing interfaces that are intuitive to the user, we would like to take the liberty of broadening the definition of affordances to include more than just design elements. In this section, we survey the literature with an eye toward the conceptual affordances of machine collaborators and discuss how they come into play in human-computer collaboration.

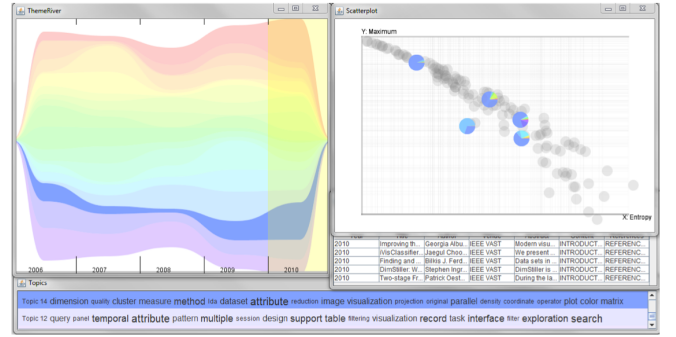
5.1 Large-scale data manipulation

As predicted by Moore’s Law [52], computational power has steadily doubled every two years for the past five and a half decades. Because of this incredible increase in processing ability, machine collaborators afford *large-scale data manipulation* at speeds and scales Fitts never could have imagined. In Visual Analytics, this computational ability has been leveraged to help analysts navigate massive datasets across many domains. For example, RP Explorer uses random projections to approximate the results of projection pursuit to find class-separating views in high-dimensional space where traditional projection pursuit can fail to converge [3]. In ParallelTopics (see Fig. 2a), computational methods for manipulating large datasets have been used to help users navigate and make sense of massive text corpora [23]. It has also been utilized to refine classification models and performing dimension reduction [18, 29, 51], interactively cluster data [4], and automatically extract transfer functions from user-selected data [57]. It has been used to suggest informative data views [83], and even to help users externalize and understand their own insight generation process [15, 39, 41, 46, 65].


5.2 Collecting and storing large amounts of data

In addition to being able to manipulate large amounts of data at incredible speed, machine collaborators are also able to *efficiently aggregate and store data* for later use. This affordance has been used to support human users in many areas where the data is being generated in large quantities and from multiple sources simultaneously. For example, systems like Verbosity [75] and others [43, 45] aggregate and store information generated by human users to create commonsense knowledge repositories. It is also used in the collection

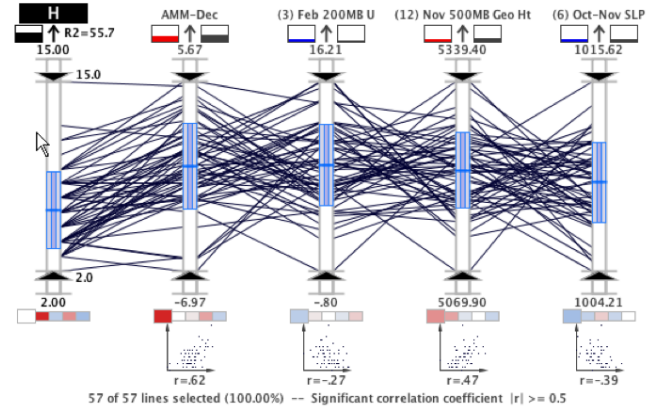
of behavioral scripts for autism treatment [11] and human-robot interaction [17], as well as collecting tags for music and image annotation [5, 12, 21, 49, 59, 62, 73, 74]. In a world that is growing ever more *big data*-centric, storage capacity and efficient retrieval are critical advantages afforded by machine collaborators.



(a) ParallelTopics [23]



(b) Saliency-Assisted Navigation [37]



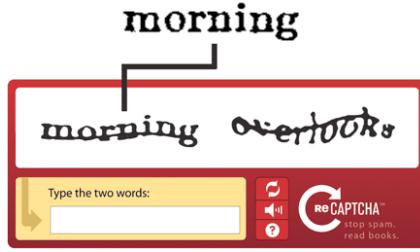
(c) MDX [66]

Figure 2: Systems leveraging machine affordances: (a) *Large-scale data manipulation*, (b) *Efficient data movement*, and (c) *Bias-free analysis*.

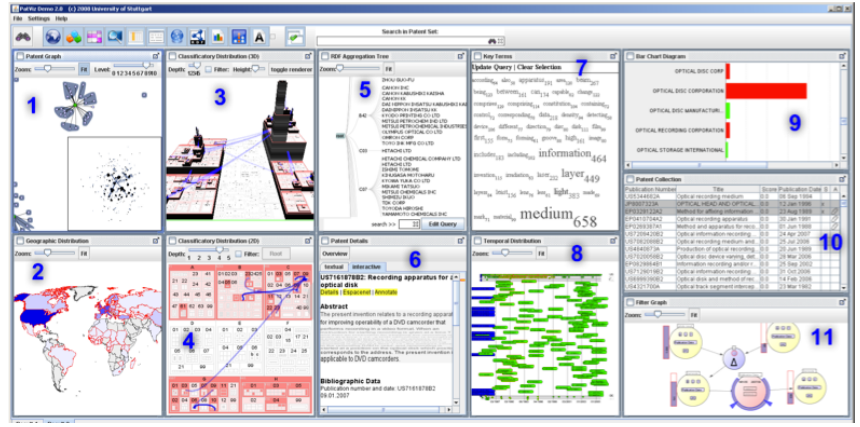
5.3 Efficient data movement

Thanks to developments in data storage, the advent of fast and reliable networking techniques and the rapid development of an always-connected society, data has been freed from its historic ties to a geographic location and machine collaborators afford very *efficient data movement*. This implies that data can be collaboratively accessed and manipulated by entities asynchronous in both time and space, with machines affording the efficient transfer of data to the right place at the right time. For example, VizWiz [10] leverages efficient data movement to connect visually-impaired users to

The Norwich line steamboat train, from New-London for Boston, this morning ran off the track seven miles north of New-London.



(a) reCAPTCHA [77]



(b) PatViz [41]

Figure 3: Systems leveraging multiple affordances: (a) reCAPTCHA [77] leverages human *visual perception* and *audiolinguistic ability* with machine *storage* and *efficient data movement* to digitize the world’s libraries. (b) PatViz leverages human *visual perception*, *visuospatial ability*, *audiolinguistic ability* and *domain knowledge* with machine *computation*, *storage* and *efficient data movement*.

sighted collaborators to get near real-time answers to visual search questions. This affordance is critical in facilitating distributed collaboration [8, 14, 19, 33, 67, 76, 80], as well as access to distributed information [33, 44, 48, 76]. Efficient data movement techniques also facilitate rapid access to data that is too large to fit in memory. This has been used to augment human visual processing using saliency modulation [37] (see Fig. 2b), as well as facilitate access to other datasets to numerous to list.

5.4 Bias-free analysis

In contrast to the human affordance of sociocultural understanding, machines afford the opportunity for *bias-free analysis*. That is, apart from human bias introduced during the programming of the system, machines are able to operate and report on numerically or computationally significant information without experiential or sociocultural influence. In Visual Analytics, we have seen this affordance leveraged to help analysts direct their attention for natural disaster prediction [66] (see Fig. 2c) as well as propose candidate visualizations for exploring high-dimensional data [68]. It has also been used to help analysts see dissimilarity to existing data-points [50], where confirmation or other bias may come into play.

6 MULTIPLE AFFORDANCES: CASE STUDIES

As stated in the introduction, while we have generally listed examples under a single main affordance, systems may utilize multiple affordances (both human and machine) in pursuit of a common goal. In this section, we analyze a few systems leveraging multiple affordances and discuss the impact of each set of design elements.

6.1 reCAPTCHA

reCAPTCHA, first introduced by Luis von Ahn et al. in [77] and later acquired by Google, is a web security mechanism that harnesses the effort of humans performing CAPTCHAs along with optical character recognition (OCR) to collaboratively digitize the world’s text corpora (see Fig. 3a).

In the first year reCAPTCHA was made available for public use, over 440 million suspicious words were correctly deciphered resulting in over 17,600 successfully transcribed books [77]. As of this writing, the system is used over 100 million times every day with an overall success rate of 96.1%, and is currently being utilized to digitize the New York Times archive as well as Google Books. Such widespread adoption and remarkable accuracy mark reCAPTCHA as one of the most widely successful human-computer collaborative initiatives to date.

We posit that the success of the reCAPTCHA system is due in part to its effective combination of human and machine affordances. After performing an initial automated recognition of a document (*computation*), suspicious or unrecognizable words are identified and transmitted (*efficient data movement*) to a collection of human collaborators for evaluation (*visual perception*) and subsequent transcription (*linguistic ability*). Through this division of labor, each party receives manageable tasks to perform according to their skills, and each set of affordances can be leveraged without overloading the collaborator.

6.2 PatViz

PatViz [41] is a Visual Analytics system for the interactive analysis of patent information (see Fig. 3b). PatViz utilizes a flexible coordinated multiple views (CMV) to support the construction of complex queries and the interactive exploration of patent result sets.

Analysis of patent information is a complex task involving the synthesis of many data dimensions. Because of this, PatViz leverages a multitude of human and machine affordances in an effort to provide intuitive views for various data types: *visual perception* for the inspection of image data contained in patent documents, *visuospatial ability* for analyzing the relationships between various patents, *audiolinguistic ability* for evaluating terminology, and *domain knowledge* for understanding the relevance of the patent to its application, as well as with machine *computation* for generating data views on the fly, *storage* for aggregating the analysts’ activity, and *efficient data movement* to provide the analyst with the appropriate information on-demand.

However, in the case of leveraging affordances, more is not always better. As articulated in the discussion of the results [41]:

One frequently expressed comment indicated that most of the patent experts never worked with a system providing interlinked and interactive visual interfaces. While this was also one of the systems properties that was most appreciated by the users, it became clear that such features are **very difficult to use** without any training.

While the machine collaborator offers many opportunities for the human to utilize many different analytical skills, it falls short in effectively leveraging these affordances by leaving the decision of when and how to select views wholly at the discretion of the human. Because so many different affordances are being leveraged, it is difficult for the human collaborators to organize their strategy in

approaching the analysis, resulting in an interface that “is difficult to comprehend... without previous instruction” [41].

7 SUGGESTED EXTENSIONS

The scope of this paper is limited to the affordances we have identified in the existing literature on human-computer collaboration and human computation; it is far from an exhaustive list of the possible affordances that exist between human and machine. We would like posit a few un- or under-explored affordances and suggest scenarios in which these affordances might prove useful.

Human Adaptability: One of the most important components of the human analytic process is the ability to take multiple perspectives on a problem, and adapt hypotheses and mental models in the wake of new information. This adaptability is critical to the successful generation of insight about large datasets. However, most work in this area has centered around *supporting* the adapting user, rather than explicitly leveraging this.

Consider the hypothetical collaborative system leveraging human adaptability along the lines suggested by Thomas and Cook [70]: as human collaborators are exploring a dataset, the system observes patterns in provenance to try to detect when an analyst has gotten “stuck” in a redundant or potentially fruitless analytical path. When this happens, the system suggests an alternative perspective or avenue for exploration. This encourages the analyst to form new hypotheses or adopt new methods of inquiry, ensuring that the analysis does not become entrenched in a local minimum.

Machine Sensing: With new developments in hardware technology rapidly becoming more readily available, there is the potential for significant advances in the kinds of sensory information that machines can make available. However, to our knowledge, this affordance has not yet been considered as part of a collaborative system.

We see potential for the utility of sensing technology as part of a human-computer collaborative team in two areas. First, sensing technology could be used to make the human collaborator aware of extrasensory information about the environment around them. Second, it could be used to respond to changes in the human collaborator themselves; for example, adapting to the user’s mental state using brain sensing technology to improve the working environment.

These represent just a brief brainstorming of potential additions to the list of affordances we have observed in the literature to date, and we hope that these ideas will inspire intellectual discourse and encourage further inquiry.

8 DISCUSSION

We close this paper with a discussion of the utility of this framework for addressing critical need in the area of human-machine collaboration, as well as its shortcomings and areas for future work.

8.1 Utility of an affordance-based framework

We claim that with the development of an affordance-based language for describing human-computer collaborative systems, we are indeed in a better position than when we first began. To validate this claim, let us return to the three questions posed in the introduction of this paper:

How do we tell if a problem would benefit from a collaborative technique? We argue that the set of problems warranting a collaborative technique is equivalent to the set problems where

there is an opportunity to effectively leverage affordances on both sides of the partnership in pursuit of the solution. By framing potential collaboration in terms of the affordances at our disposal, we can then consider which of these affordances could be used to approach a problem and construct a solution.

How do we decide which tasks to delegate to which party, and when? In adopting this language, we are deliberately moving away from terminology that encourages us to speak in terms of deficiencies; that is, we need the human because computers are bad at X, etc. Instead of deciding who gets (stuck with) which task, we begin to reason about who can contribute to the collective goal at each stage. The answer may not be *only the human*, or *only the machine*, but could in fact be *both*. By designing such that all parties are aware of the affordances made available to them by their collaborators, we encourage the development of more flexible procedures for collective problem-solving.

How does one system compare to others trying to solve the same problem? Of the contributions made by this framework, we believe that providing a common language for discussing human-computer collaborative systems is its greatest strength. We are able to talk about which affordances are being leveraged, and use these to compare and contrast between systems. We may also be able to make hypotheses about how these choices of affordances influence the resulting solutions by comparing performance measures. However, this language does not yet afford a robust, theoretical comparison. To achieve this, we must first build our understanding of the mechanisms underlying these affordances and their associated costs.

8.2 Complexity measures for Visual Analytics

While we believe that this framework provides an important foundation for developing a common language, it is only the first of many steps toward a rich vocabulary for describing human-computer collaborative systems. Consider for example the plethora of human computation systems for image labeling that we have reviewed in this work: the ESP Game [73], Ka-captcha [21], KissKissBan [33], LabelMe [59] and Phetch [74]. Each system leverages the visual perception and linguistic abilities of the human users, and the aggregative capacity of the machine. Given that these systems are all addressing very similar problems using a similar approach, how do they compare to one another? We argue that is it critical to develop a common language not just for describing which affordances are being leveraged, but *how much* and *how well*.

The National Science Foundation CISE directorate has called for the development of theoretical measures for systems involving human computation, calling this one of the five most important questions facing computer science today [79]. This need was reiterated at the CHI2011 workshop on Crowdsourcing and Human Computation [42]. Can we begin to describe the complexity of human-computer collaborative systems with a robust language parallel to describing the complexity of an algorithmic system?

Researchers in the field of Artificial Intelligence have begun to imagine the concept of complexity measures for systems involving human contribution. Shahaf and Amir define a *Human-Assisted Turing Machine* using the human as an oracle with known complexity [61]. In this work, they demonstrate that much of the standard theoretical language holds true, including *algorithmic complexity*, *problem complexity*, *complexity classes* and more. However, they also raise several questions that remain unanswered:

- First, **what is the best way to measure human work?** In terms of *human time*, *space*, or *utility*? Should we consider the *input size*, that is, how much data does the human need to process? Or to compensate for compression, should we be measuring *information density* instead?

- Second, **how can we assess this human work in practice?** Through *empirical evaluation* of a sample population's performance on a given task, we can begin to understand how the average human performs, but this information is task-specific. Perhaps more broadly applicable would be to develop a set of *canonical actions* that humans can perform with known complexity, but compiling this list is nontrivial.
- Finally, **how do we account for individual differences in human operators?** Perhaps the problem under consideration utilizes skills or knowledge not common to every user (such as bilingual translation). In this case, a general model of humans is insufficient; instead, we need to understand the complexity of the individual candidate. This requires the development of algorithmic systems that are to be able to effectively and efficiently utilize the affordances provided by the humans available to them, rather than only the optimal human collaborator under perfect conditions.

These areas afford many rich opportunities for collaboration with our colleagues in theoretical computer science, as well as in psychology and neuroscience. By engaging in the interdisciplinary pursuit of answers around human affordances, we hope to construct a more complete picture of insight generation, the mechanisms of human understanding, and the the analytic process as a whole.

9 CONCLUSION

In this article, we have presented a comprehensive survey of work in the area of human-computer collaborative systems and human computation. Out of 1,271 papers reviewed, we selected a corpus of 49 publications that represent the state of the art, and from these papers we identify patterns of human- and machine-intelligence affordances. These affordances form the basis of a common framework for understanding and discussing this collection of work. We argue for the utility of this framework, then discuss some of its shortcomings and identify unexplored avenues for extending this line of inquiry.

While we have concentrated our efforts on systems explicitly labeled as human-computer collaboration, mixed-initiative, or human computation, we posit that the framework presented here will benefit the field of Visual Analytics as a whole. While there has been remarkable progress in the development of novel solutions to support analytic processes, we have not yet fully realized our potential as a systematic science that builds and organizes knowledge in the form of testable theories and predictions. In presenting a preliminary framework for describing and comparing systems involving human and machine collaborators, we aspire to lay the foundation for a more rigorous analysis of the tools and approaches presented by our field, thereby paving the way for the construction of an increasingly robust understanding of analytical reasoning and how to best support insight generation.

REFERENCES

- [1] E. Adar. Why i hate mechanical turk research (and workshops). In *Proc. of the 29th SIGCHI Conf. on Human factors in computing systems: Workshop on Crowdsourcing and Human Computation*. ACM, 2011.
- [2] T. Amabile. *Creativity in context: Update to "the social psychology of creativity"*. Westview press, 1996.
- [3] A. Anand, L. Wilkinson, and T. Dang. Using random projections to identify class-separating variables in high-dimensional spaces. In *Visual Analytics Science and Technology (VAST), Conf. on*, pages 263–264. IEEE, 2011.
- [4] G. Andrienko, N. Andrienko, S. Rinzivillo, M. Nanni, D. Pedreschi, and F. Giannotti. Interactive visual clustering of large collections of trajectories. In *Visual Analytics Science and Technology (VAST) 2009, Symposium on*, pages 3–10. IEEE, 2009.
- [5] L. Barrington, D. O'Malley, D. Turnbull, and G. Lanckriet. User-centered design of a social game to tag music. In *Proc. of the ACM SIGKDD Workshop on Human Computation*, pages 7–10. ACM, 2009.
- [6] M. Bell, S. Reeves, B. Brown, S. Sherwood, D. MacMillan, J. Ferguson, and M. Chalmers. Eyespy: supporting navigation through play. In *Proc. of the 27th SIGCHI Conf. on Human factors in computing systems*, pages 123–132. ACM, 2009.
- [7] B. Berger and T. Leighton. Protein folding in the hydrophobic-hydrophilic (hp) model is np-complete. *Journal of Computational Biology*, 5(1):27–40, 1998.
- [8] M. Bernstein, G. Little, R. Miller, B. Hartmann, M. Ackerman, D. Karger, D. Crowell, and K. Panovich. Soylent: a word processor with a crowd inside. In *Proc. of the 23rd symposium on User interface software and technology*, pages 313–322. ACM, 2010.
- [9] E. Bertini and D. Lalanne. Investigating and reflecting on the integration of automatic data analysis and visualization in knowledge discovery. *ACM SIGKDD Explorations Newsletter*, 11(2):9–18, 2010.
- [10] J. Bigham, C. Jayant, H. Ji, G. Little, A. Miller, R. Miller, R. Miller, A. Tatarowicz, B. White, S. White, et al. Vizwiz: nearly real-time answers to visual questions. In *Proc. of the 23rd symposium on User interface software and technology*, pages 333–342. ACM, 2010.
- [11] F. Boujarwah, J. Kim, G. Abowd, and R. Arriaga. Developing scripts to teach social skills: Can the crowd assist the author? In *Workshops at the Twenty-Fifth AAAI Conf. on Artificial Intelligence*, 2011.
- [12] I. Bowman, S. Joshi, and J. Van Horn. Query-based coordinated multiple views with feature similarity space for visual analysis of mri repositories. In *Visual Analytics Science and Technology (VAST), Conf. on*, pages 267–268. IEEE, 2011.
- [13] M. Chen, D. Ebert, H. Hagen, R. Laramée, R. Van Liere, K. Ma, W. Ribarsky, G. Scheuermann, and D. Silver. Data, information, and knowledge in visualization. *Computer Graphics and Applications, IEEE*, 29(1):12–19, 2009.
- [14] Y. Chen, J. Alsakran, S. Barlowe, J. Yang, and Y. Zhao. Supporting effective common ground construction in asynchronous collaborative visual analytics. In *Visual Analytics Science and Technology (VAST), Conf. on*, pages 101–110. IEEE, 2011.
- [15] Y. Chen, S. Barlowe, and J. Yang. Click2annotate: Automated insight externalization with rich semantics. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 155–162. IEEE, 2010.
- [16] Y. Chen, B. Liem, and H. Zhang. An iterative dual pathway structure for speech-to-text transcription. In *Human Computation: Papers from the AAAI Workshop (WS-11-11)*. San Francisco, CA, 2011.
- [17] S. Chernova, J. Orkin, and C. Breazeal. Crowdsourcing hri through online multiplayer games. In *Proc. Dialog with Robots: AAAI fall symposium*, 2010.
- [18] J. Choo, H. Lee, J. Kihm, and H. Park. ivisclassifier: An interactive visual analytics system for classification based on supervised dimension reduction. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 27–34. IEEE, 2010.
- [19] K. Chung, C. Chiu, X. Xiao, and P. Chi. Stress outsourced: a haptic social network via crowdsourcing. In *Proc. of the 27th SIGCHI Conf. on Human factors in computing systems*, pages 2439–2448. ACM, 2009.
- [20] S. Cooper, F. Khatib, A. Treuille, J. Barbero, J. Lee, M. Beenen, A. Leaver-Fay, D. Baker, Z. Popovic, et al. Predicting protein structures with a multiplayer online game. *Nature*, 466(7307):756–760, 2010.
- [21] B. Da Silva and A. Garcia. Ka-captcha: An opportunity for knowledge acquisition on the web. In *Proc. of the national Conf. on Artificial Intelligence*, volume 22, page 1322, 2007.
- [22] S. Dekker and D. Woods. Maba-maba or abracadabra? progress on human-automation coordination. *Cognition, Technology & Work*, 4(4):240–244, 2002.
- [23] W. Dou, X. Wang, R. Chang, and W. Ribarsky. Paralleltopics: A probabilistic approach to exploring document collections. *Visual Analytics Science and Technology (VAST), Conf. on*, 2011.
- [24] P. Dourish. *Where the action is: the foundations of embodied interaction*. The MIT Press, 2004.
- [25] H. Fastl and E. Zwicker. *Psychoacoustics: facts and models*, volume 22. Springer-Verlag New York Inc, 2007.
- [26] G. Ferguson and J. Allen. Trips: An integrated intelligent problem-solving assistant. In *Proc. of the National Conf. on Artificial Intelligence*, pages 567–573. JOHN WILEY & SONS LTD, 1998.
- [27] P. Fitts. Human engineering for an effective air-navigation and traffic-control system. 1951.
- [28] R. Fuchs, J. Waser, and M. Groller. Visual human+ machine learning. *Visualization and Computer Graphics, Transactions on*, 15(6):1327–1334, 2009.
- [29] S. Garg, I. Ramakrishnan, and K. Mueller. A visual analytics approach to model learning. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 67–74. IEEE, 2010.
- [30] W. Gaver. Technology affordances. In *Proc. of the SIGCHI Conf. on Human factors in computing systems: Reaching through technology*, pages 79–84. ACM, 1991.
- [31] J. Gibson. The theory of affordances. *Perceiving, acting, and knowing*, pages

	Human Affordances						Machine Affordances			
	Visual	Spatial	Aud/Ling.	Creativity	Social	Domain	Comp.	Storage	Moving	Bias-Free
PatViz [41]	✓	✓	✓			✓	✓	✓	✓	
CrowdSearch [80]	✓						✓		✓	
ParallelTopics [23]	✓		✓				✓			
Dissimilarity [50]	✓					✓	✓			✓
VH+ML [28]	✓						✓			
Implicit tagging [62]	✓						✓			
reCAPTCHA [77]	✓		✓					✓	✓	
VizWiz [10]	✓				✓				✓	
Phetch [74]	✓		✓	✓	✓			✓		
ESP Game [73]	✓		✓	✓				✓		
KissKissBan [33]	✓		✓	✓				✓		
LabelMe [59]	✓		✓	✓				✓		
Ka-captcha [21]	✓		✓	✓				✓		
PeekABoom [76]	✓							✓		
MRI [12]		✓				✓	✓	✓	✓	
iView [83]		✓				✓	✓		✓	
iVisClassifier [18]		✓				✓	✓		✓	
Saliency [37]		✓				✓	✓		✓	
RP Explorer [3]		✓					✓		✓	✓
DimStiller [36]		✓				✓	✓	✓		
WireVis [46]		✓				✓	✓	✓		
Action trails [65]		✓				✓	✓	✓		
NetClinic [47]		✓				✓	✓			
Trajectories [4]		✓				✓	✓			
Risk assessment [51]		✓				✓	✓			
Automatic transfer functions [57]		✓				✓	✓			✓
MDX [66]		✓				✓	✓			✓
Automated+viz [68]		✓				✓	✓			✓
CzSaw [39]		✓					✓	✓		
Fold.it [20]		✓					✓			
HRI scripts [17]		✓			✓			✓		
Animated agents for VR [56]		✓			✓			✓		
VA Model-learning [29]		✓				✓				
EyeSpy [6]		✓						✓		
MonoTrans2 [35]			✓				✓	✓		
CastingWords [16]			✓				✓			
Click2Annotate [15]						✓	✓	✓		
Wrangler [40]						✓	✓			
Soylent [8]			✓						✓	
Crowdsourced solutions [67]				✓					✓	
Crowdsourced design [81]				✓					✓	
Stress OutSourced [19]					✓				✓	
PageHunt [48]			✓	✓	✓			✓		
Herd It [5]			✓	✓				✓		
TagATune [44]			✓	✓				✓		
MajorMiner [49]			✓	✓				✓		
Autism scripts [11]					✓			✓		
Social Games [43]					✓			✓		
Common Consensus [45]					✓			✓		
Verbosity [75]					✓			✓		

Table 1: A table of all surveyed human-computer collaborative systems and the affordances they leverage. The human affordances listed are (in order): *visual perception*, *visuospatial ability*, *audiolinguistic ability*, *creativity*, *sociocultural awareness*, and *domain knowledge*. The machine affordances listed are (in order): *large-scale data manipulation*, *collecting and storing large amounts of data*, *efficient data movement*, and *bias-free analysis*.

- 67–82, 1977.
- [32] J. Gibson. *The ecological approach to visual perception*. Lawrence Erlbaum, 1986.
 - [33] C. Ho, T. Chang, J. Lee, J. Hsu, and K. Chen. Kisskissban: a competitive human computation game for image annotation. In *Proc. of the SIGKDD Workshop on Human Computation*, pages 11–14. ACM, 2009.
 - [34] E. Horvitz. Principles of mixed-initiative user interfaces. In *Proc. of the SIGCHI Conf. on Human factors in computing systems: the CHI is the limit*, pages 159–166. ACM, 1999.
 - [35] C. Hu, B. Bederson, P. Resnik, and Y. Kronrod. Monotrans2: A new human computation system to support monolingual translation. In *Proc. of the 29th SIGCHI Conf. on Human factors in computing systems*, pages 1133–1136. ACM, 2011.
 - [36] S. Ingram, T. Munzner, V. Irvine, M. Tory, S. Bergner, and T. Miller. Dimstiller: Workflows for dimensional analysis and reduction. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 3–10. IEEE, 2010.
 - [37] C. Ip and A. Varshney. Saliency-assisted navigation of very large landscape images. *Visualization and Computer Graphics, Transactions on*, 17(12):1737–1746, 2011.
 - [38] N. Jordan. Allocation of functions between man and machines in automated systems. *Journal of applied psychology*, 47(3):161, 1963.
 - [39] N. Kadivar, V. Chen, D. Dunsmuir, E. Lee, C. Qian, J. Dill, C. Shaw, and R. Woodbury. Capturing and supporting the analysis process. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 131–138. IEEE, 2009.
 - [40] S. Kandel, A. Paepcke, J. Hellerstein, and J. Heer. Wrangler: Interactive visual specification of data transformation scripts. In *Proc. of the 2011 Conf. on Human factors in computing systems*, pages 3363–3372. ACM, 2011.
 - [41] S. Koch, H. Bosch, M. Giereth, and T. Ertl. Iterative integration of visual insights during patent search and analysis. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 203–210. IEEE, 2009.
 - [42] A. Kulkarni. The complexity of crowdsourcing: Theoretical problems in human computation. In *CHI Workshop on Crowdsourcing and Human Computation*, 2011.
 - [43] Y. Kuo, J. Lee, K. Chiang, R. Wang, E. Shen, C. Chan, and J. Hsu. Community-based game design: experiments on social games for commonsense data collection. In *Proc. of the ACM SIGKDD Workshop on Human Computation*, pages 15–22. ACM, 2009.
 - [44] E. Law, L. Von Ahn, R. Dannenberg, and M. Crawford. Tagatune: A game for music and sound annotation. *Proc. of ISMIR (Vienna, Austria)*, 2007.
 - [45] H. Lieberman, D. Smith, and A. Teeters. Common consensus: a web-based game for collecting commonsense goals. In *ACM Workshop on Common Sense for Intelligent Interfaces*, 2007.
 - [46] H. Lipford, F. Stukes, W. Dou, M. Hawkins, and R. Chang. Helping users recall their reasoning process. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 187–194. IEEE, 2010.
 - [47] Z. Liu, B. Lee, S. Kandula, and R. Mahajan. Netclinic: Interactive visualization to enhance automated fault diagnosis in enterprise networks. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 131–138. IEEE, 2010.
 - [48] H. Ma, R. Chandrasekar, C. Quirk, and A. Gupta. Page hunt: using human computation games to improve web search. In *Proc. of the ACM SIGKDD Workshop on Human Computation*, pages 27–28. ACM, 2009.
 - [49] M. Mandel and D. Ellis. A web-based game for collecting music metadata. *Journal of New Music Research*, 37(2):151–165, 2008.
 - [50] M. Migut, J. van Gemert, and M. Worring. Interactive decision making using dissimilarity to visually represented prototypes. In *Visual Analytics Science and Technology (VAST), Conf. on*, pages 141–149. IEEE, 2011.
 - [51] M. Migut and M. Worring. Visual exploration of classification models for risk assessment. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 11–18. IEEE, 2010.
 - [52] G. Moore et al. Cramming more components onto integrated circuits. *Proc. of the IEEE*, 86(1):82–85, 1998.
 - [53] D. Norman. *The design of everyday things*. Basic books, 2002.
 - [54] H. Price. The allocation of functions in systems. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 27(1):33–45, 1985.
 - [55] A. Quinn and B. Bederson. Human computation: a survey and taxonomy of a growing field. In *Proc. of the 29th SIGCHI Conf. on Human factors in computing systems*, pages 1403–1412. ACM, 2011.
 - [56] J. Rickel and W. Johnson. Animated agents for procedural training in virtual reality: Perception, cognition, and motor control. *Applied Artificial Intelligence*, 13(4-5):343–382, 1999.
 - [57] M. Ruiz, A. Bardera, I. Boada, and I. Viola. Automatic transfer functions based on informational divergence. *Visualization and Computer Graphics, Transactions on*, 17(12):1932–1941, 2011.
 - [58] M. Runco. *Creativity: Theories and themes: Research, development, and practice*. Academic Press, 2007.
 - [59] B. Russell, A. Torralba, K. Murphy, and W. Freeman. Labelme: a database and web-based tool for image annotation. *Int'l journal of computer vision*, 77(1):157–173, 2008.
 - [60] P. Shah and A. Miyake. *The Cambridge handbook of visuospatial thinking*. Cambridge Univ Pr, 2005.
 - [61] D. Shahaf and E. Amir. Towards a theory of ai completeness. In *8th International symposium on logic formalizations of commonsense reasoning*, 2007.
 - [62] P. Shenoy and D. Tan. Human-aided computing: Utilizing implicit human processing to classify images. In *Proc. of the 26th SIGCHI Conf. on Human factors in computing systems*, pages 845–854. ACM, 2008.
 - [63] T. Sheridan. Function allocation: algorithm, alchemy or apostasy? *Int'l Journal of Human-Computer Studies*, 52(2):203–216, 2000.
 - [64] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Visual Languages, 1996. Proc., IEEE Symposium on*, pages 336–343. IEEE, 1996.
 - [65] Y. Shrinivasan, D. Gotzy, and J. Lu. Connecting the dots in visual analysis. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 123–130. IEEE, 2009.
 - [66] C. Steed, J. Swan, T. Jankun-Kelly, and P. Fitzpatrick. Guided analysis of hurricane trends using statistical processes integrated with interactive parallel coordinates. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 19–26. IEEE, 2009.
 - [67] Y. Tanaka, Y. Sakamoto, and T. Kusumi. Conceptual combination versus critical combination: Devising creative solutions using the sequential application of crowds. In *Proc. of the 33rd Conf. of the Cognitive Science Society*, 2011.
 - [68] A. Tatu, G. Albuquerque, M. Eisemann, J. Schneidewind, H. Theisel, M. Magnork, and D. Keim. Combining automated analysis and visualization techniques for effective exploration of high-dimensional data. In *Visual Analytics Science and Technology (VAST), Symposium on*, pages 59–66. IEEE, 2009.
 - [69] L. Terveen. Overview of human-computer collaboration. *Knowledge-Based Systems*, 8(2-3):67–81, 1995.
 - [70] J. Thomas and K. Cook. *Illuminating the path: The research and development agenda for visual analytics*, volume 54. IEEE, 2005.
 - [71] R. Valdés-Pérez. Principles of human-computer collaboration for knowledge discovery in science. *Artificial Intelligence*, 107(2):335–346, 1999.
 - [72] L. Von Ahn. *Human Computation*. PhD thesis, Carnegie Mellon University, 2005.
 - [73] L. Von Ahn and L. Dabbish. Labeling images with a computer game. In *Proc. of the 22nd SIGCHI Conf. on Human factors in computing systems*, pages 319–326. ACM, 2004.
 - [74] L. Von Ahn, S. Ginosar, M. Kedia, R. Liu, and M. Blum. Improving accessibility of the web with a computer game. In *Proc. of the 24th SIGCHI Conf. on Human factors in computing systems*, pages 79–82. ACM, 2006.
 - [75] L. Von Ahn, M. Kedia, and M. Blum. Verbosity: a game for collecting commonsense facts. In *Proc. of the 24th SIGCHI Conf. on Human factors in computing systems*, pages 75–78. ACM, 2006.
 - [76] L. Von Ahn, R. Liu, and M. Blum. Peekaboom: a game for locating objects in images. In *Proc. of the 24th SIGCHI Conf. on Human factors in computing systems*, pages 55–64. ACM, 2006.
 - [77] L. Von Ahn, B. Maurer, C. McMillen, D. Abraham, and M. Blum. recaptcha: Human-based character recognition via web security measures. *Science*, 321(5895):1465–1468, 2008.
 - [78] L. Vygotsky. Thought and word. 1962.
 - [79] J. Wing. Five deep questions in computing. *Communications of the ACM*, 51(1):58–60, 2008.
 - [80] T. Yan, V. Kumar, and D. Ganesan. Crowdsearch: exploiting crowds for accurate real-time image search on mobile phones. In *Proc. of the 8th Int'l Conf. on Mobile systems, applications, and services*, pages 77–90. ACM, 2010.
 - [81] L. Yu and J. Nickerson. Cooks or cobblers?: crowd creativity through combination. In *Proc. of the 29th SIGCHI Conf. on Human factors in computing systems*, pages 1393–1402. ACM, 2011.
 - [82] M. Yuen, L. Chen, and I. King. A survey of human computation systems. In *Computational Science and Engineering, Int'l Conf. on*, volume 4, pages 723–728. IEEE, 2009.
 - [83] Z. Zheng, N. Ahmed, and K. Mueller. iVIEW: A feature clustering framework for suggesting informative views in volume visualization. *Visualization and Computer Graphics, Transactions on*, 17(12):1959–1968, 2011.