Baloo: Personal Informatics for Decision Making

Edward K. Smith
University of Maryland, College Park
tedks@cs.umd.edu

James Purtilo
University of Maryland, College Park
purtilo@cs.umd.edu

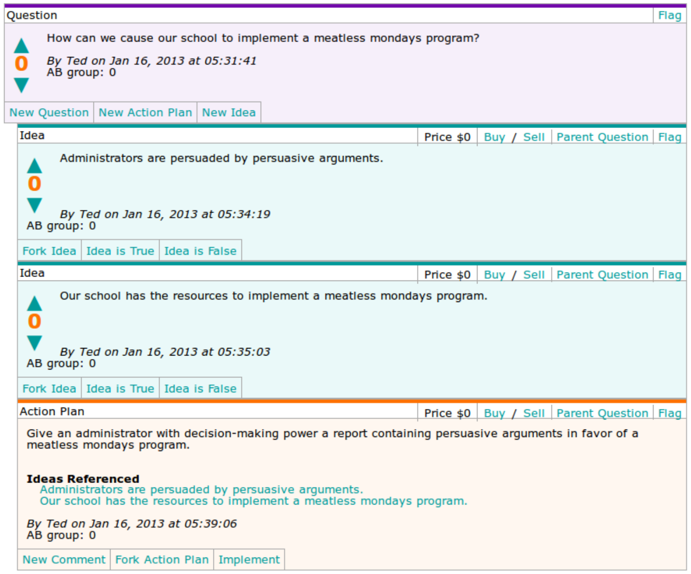

Figure 1: A screenshot of Baloo.

Abstract
While decision making is a well-studied area in cognitive psychology, few quantitative tools exist to aid decision making, mitigate the array of biases present in humans, and improve calibration. We are developing such a system, Baloo, that assists users in quantifying and tracking their belief in factual propositions, with the goal of combining these propositions into successful plans. With Baloo, individuals and groups can easily assess the relationship between various beliefs, and improve their predictive accuracy over time. In this paper, we describe the design of Baloo.

ACM Classification Keywords
H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous

Introduction
Decision making has been the focus of active research in cognitive psychology for decades. Cognitive psychologists have extensive understanding of the heuristics that humans use to assess likelihood and form plans. These heuristics have reliable failure modes, or *biases*. While these biases are well-known to cognitive psychologists, they are mostly unknown to the general public. As a result, most decision making is reliably biased.

For example, any decision is based on the expectation of future events. However, humans reliably violate the laws of probability in their assessment of the likelihood of events, often assigning greater probability to $P(A\&B)$ than to $P(A)$, often referred to as the conjunction fallacy [8]. Even on predictions of single propositions, even experts are reliably overconfident [7].

With training, humans can become better *calibrated*: their confidence in a statement can be predictive of its truth. There are some systems which attempt to train users' calibration [2, 1], but these systems are ill-suited for decision making because they do not track relationships between predictions or allow multi-level analysis. Baloo allows users to track predictions and improve calibration, but also rewards users for accuracy, and structures information to facilitate decision making.

We have developed a prototype system, Baloo, which attempts to use this information to mitigate bias in decision making. Baloo allows individual or group users to quantify their belief in factual propositions or plans, aggregates this belief when appropriate, and presents this information its user. As users continue to use Baloo, they receive feedback on their predictive accuracy, and can use this feedback to improve their calibration.

An example of the system is available for public use at http://baloo.umd.edu. We intend to release Baloo's source code under the terms of the GNU Affero General Public License. This version of Baloo was designed and implemented by the Fall 2012 semester of UMD's Software Engineering class.

Baloo Design

Baloo is a spiritual successor to MMOWGLI [5], a similar system developed by the US Navy and Institute for the Future[1] to generate possible solutions to difficult open problems. MMOWGLI is phrased in terms of a game where the objective is to obtain points through submitting insightful *ideas* and *plans*, with the eventual goal of collaborative producing a solution to large-scale problems. Baloo retains these classes of interactions, but adds additional structure to support building specific multi-level plans. Baloo also differs from MMOWGLI by allowing users to quantify their belief in ideas and plans, and by propagating this information between related interactions.

Baloo contains three types of interactions: ideas, plans, and problems. In Baloo, an *idea* is a prediction about the state of the world. An idea includes a testable statement, a probability estimation, and a deadline, at which point the truth of the statement will be known. A *plan* is a sequence of actions designed to bring about a given effect. Ideas and plans can be *linked*, indicating that the truth of Idea B depends on the truth of Idea A. These interactions occur in the context of a *problem*. Any problem can itself contain child problems representing subsets of the larger decision. The goal of any plan in Baloo is to solve its parent problem. When the truth of an idea or the success of a plan is known, these interactions can be *closed* with this information.

Baloo collects the estimated probability of any idea being true, or any plan succeeding, from its users. With this information, and the dependency information generated by linking, Baloo can set upper bounds on the probabilities of ideas and plans. This mechanism is designed to inoculate against the conjunction fallacy.

These interactions are displayed in the form of a hierarchical tree (see figure 1). Links between interactions

[1] http://www.iftf.org/

are currently listed in the body of the interactions.

Baloo in the Personal Informatics Stage Model

In this section, we outline Baloo's structure in the stage model [6] of personal informatics, using the example of an activist group as the user.

Preparation Currently, before using Baloo, any individual user must consider the decision they plan to make, and decide if it should be viewed by groups, or by them alone. If the decision is relevant to another decision, they may choose to embed the problem as a subproblem in an existing problem, to take advantage of the existing ideas and plans tracked in Baloo.

Collection All collection in Baloo is manual, and comes in the form of the user creating problems, ideas, and plans within Baloo. This form of collection is inspired both by Baloo's predecessor, MMOWGLI, and by mind-mapping software that displays relationships between ideas.

The core information Baloo collects is a probability from users representing the likelihood an idea is true, or the chance a plan will succeed. In individual cases, this can only be collected directly, as a numeric probability.

However, in group use, Baloo aggregates belief using *prediction markets*. Prediction markets are speculative markets in which users buy and sell shares in the truth of a prediction; these markets have been shown to be highly accurate in predicting events such as elections [3].

Baloo collects dependencies between interactions when they are created, by allowing the user to select ideas the new idea or plan depends on from a list of existing ideas within the scope of the new idea or plan. After the truth of an idea or the success of a plan is known, users can *close* the interaction as true or false. Closing an interaction liquidates its prediction markets, and updates individual accuracy records.

In our hypothetical activist group, a card chain such as Figure 1 might be created, containing ideas related to the willingness of an administration to engage with the activist group, and plans corresponding with various strategies. Over the course of the campaign, the activists revise their beliefs by trading within the prediction markets.

Integration Integration in Baloo is entirely system-driven. Given the user's probability estimate and dependency information for all ideas and plans, Baloo can automatically make inferences about various other properties of various ideas or plans. For example, the upper bound of any interaction's probability is the maximum of the probability of any interaction on which it depends ($P(A\&B)$ is at most $P(A)$). Additionally, Baloo can refer to past predictions of specific users to determine the actual accuracy of predictions. For example, if an activist predicts with 70% confidence that an administrator will meet with them, and that activist's predictions with similar confidence have been true 60% of the time, Baloo can internally represent the probability of the idea as 60%.

Reflection and Action In group contexts, Baloo converts prediction market prices to probabilities and displays those probabilities. Baloo sorts the list of ideas and plans in order of probability, allowing rapid assessment of potentially successful plans. Each user can also view their past predictions, and their current calibration.

After the facts of particular situations become known, and interactions are closed, their prediction markets will economically reward users who correctly predicted the outcome. Repeated erroneous prediction will result in less ability to influence future decisions. In this sense, a group using Baloo to determine probable successful plans is implementing a form of futarchy: governance via prediction [4]. A user's capital accumulation in Baloo also provides a crude but effective accuracy metric.

Future Work

We have built a prototype version of Baloo that supports group use, prediction markets, and rudimentary dependency tracking, but lacks many features. We envision a wide variety of extensions that could make Baloo more useful, and hope to solicit further enhancements from workshop participants.

One current weakness of Baloo is the lack of *conditional* probability knowledge and reflection. We hope to extend Baloo to collect and integrate conditional probabilities of ideas and plans, to answer questions such as "If Idea A is true, how likely is it that Plan B will succeed?" We are interested in exploring multiple avenues of collecting and reflecting this data to users, integrating existing work in visualizing conditional probability.

People tend to ask for advice when making difficult decisions. Facilitating advice while retaining privacy is a logical next step for Baloo. Baloo would enable decision makers to quantitatively aggregate advice, while hiding the identity of the advisers when desired.

Finally, in its current form, Baloo lacks strong visualizations of the information it contains, primarily because none of its authors (including the authors of this paper) are particularly skilled in information visualization. We hope to enhance Baloo with informative visualizations, including conditional probabilities, impact of various ideas or plans, and relationships between ideas and plans.

We are currently completing the single-player mode of Baloo, and refining the group mode of Baloo with enhanced visualizations of probability distributions and dependencies. When this is complete, we hope to conduct usability studies to assess the degree to which users are able to understand Baloo's debiasing systems, and use them to improve calibration.

References

[1] Intrade: An online prediction market. http://intrade.com.

[2] Predictionbook: An online prediction-recording system. http://predictionbook.com.

[3] Berg, J., and Rietz, T. Prediction markets as decision support systems. *Information Systems Frontiers 5*, 1 (2003), 79–93.

[4] Hanson, R. Shall we vote on values, but bet on beliefs? *Working Paper, George Mason University* (2003).

[5] Jenson, G., and Tester, J. Government for the 100%: Using games to democratize innovation and innovate democracy, 2012.

[6] Li, I., Dey, A., and Forlizzi, J. A stage-based model of personal informatics systems. In *Proceedings of the 28th international conference on Human factors in computing systems*, ACM (2010), 557–566.

[7] Lichtenstein, S., Fischhoff, B., and Phillips, L. Calibration of probabilities: The state of the art to 1980. Tech. rep., DTIC Document, 1981.

[8] Tversky, A., and Kahneman, D. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological review 90*, 4 (1983), 293.