Targeted Trust-Based Merging of Customers' Opinions

Jurij Ružejnikov[1,2(✉)] and Tatiana Valentine Guy[2]

[1] Czech University of Life Sciences Prague, Kamýcká 129, 165 00 Praha - Suchdol, Czech Republic
{ruzejnikov,guy}@utia.cas.cz

[2] Institute of Information Theory and Automation, Czech Academy of Sciences, Pod Vodárenskou věží 4, 182 08 Praha 8, Czech Republic

Abstract. In this article, we investigate how a rational agent forms their opinion based on prior knowledge, available information, and the opinions of other agents. We propose methodology of how to purposefully merge agent's opinion and expert opinions. We describe the agent's opinion and the opinions of experts in the form of distributions. Formulating opinion formation as a decision-making task and solve it using Fully Probabilistic Design (FPD). To demonstrate our approach, we apply the solution on simulated data describing features of mobile phone brands. Methodology is verified on a test bed example of choosing a mobile phone brand based on expert opinions while taking into account agent's trust in experts.

Keywords: Decision-making · Fully probabilistic design · Knowledge representation · Opinion dynamics · Opinion merging · Trust

1 Introduction

Opinion dynamics is a compelling field that explores how individual beliefs and preferences evolve through interactions and information exchange. Among the various methods employed to model these dynamics, linear pooling of opinions [1] stands out as a widely used technique. This method aggregates expert opinions into a single, coherent viewpoint by calculating a weighted sum of their inputs. Similarly, logarithmic opinion pooling [3] employs the weighted geometric mean to synthesize expert insights, providing a robust framework for consensus modelling.

Beyond these foundational approaches, innovative strategies such as consensual modelling [5] have emerged, leveraging social trust networks to enhance group decision-making processes. For example, in [14], experts collaborate to form a consensus, akin to an individual seeking diverse opinions from friends before making a significant purchase, such as a new computer.

Another compelling method, merging opinions by social sampling of posteriors [13], introduces a dynamic way to update opinions based on external information sources, paralleling the challenges we address in this paper. Additionally,

W. Senn et al. (Eds.): ICANN 2025, LNCS 16071, pp. 351–365, 2026.
https://doi.org/10.1007/978-3-032-04555-3_29

Bayesian updates can refine a modeller's opinion by integrating expert data [9], a concept we will explore further. A notable example is the Bayesian hierarchical model [11], which uses Kullback-Leibler divergence to objectively determine the influence of each expert and produce a calibrated consensus from multiple sources. Our approach, however, addresses the challenge from a distinct, single-agent perspective. Instead of seeking an objective consensus, we focus on how a rational agent personally forms an opinion by merging their own prior knowledge with the subjective trust they place in external experts. We explore how this agent can purposefully integrate diverse opinions, such as a consumer leveraging friend recommendations or online reviews before purchasing a new mobile phone. This paper develops a framework based on divergence minimization [2] and Fully Probabilistic Design [8] to model this subjective process and enhance the merging of opinion-defining distributions, ultimately aimed at improving decision-making outcomes.

The variety of proposed solutions highlights that the problem is far from resolved. In our approach, we emphasize that the data represent *subjective* opinions rather than objective truths. This distinction is essential, as it enables a nuanced update process that integrates expert insights while also reflecting the modeller's trust in these opinions and their own level of certainty. We quantify opinions as parametrised distributions of a random variable that encapsulates the subject at hand.

To showcase the effectiveness of our method, we will use the selection of a mobile phone brand as a practical case study, drawing on data collected from real participants. This scenario not only demonstrates the application of our approach but also underscores the complexities involved in merging diverse opinions in a real-world context. By addressing these challenges, we aim to provide valuable insights into the decision-making process and contribute to the ongoing conversation in opinion dynamics. The **contributions of the paper** are as follows:

- The paper introduces the innovative framework for dynamical merging of opinions that effectively differentiates between subjective insights and objective data. It allows for a nuanced integration of expert opinions while respecting the targeted scenario that involves multi-attribute preferences and expert prior opinions. This proposed approach not only enhances decision-making but also empowers users to make *informed choices* based on understanding of expert insights.
- The paper explores the role of the modeller's trust in expert opinions as well as their own level of certainty, emphasizing how these factors significantly impact the opinion merging process and shape decision outcomes. The research highlights the important relationship between trust and certainty, offering clear insights that can improve decision-making processes.
- Through a series of experiments conducted on real participant data, it validates the proposed method and demonstrates its effectiveness. This empirical validation not only demonstrates the effectiveness of the framework but also showcases its practical applicability in real-world scenarios.

Here are several potential use cases that showcase the relevance of the proposed framework for *dynamical* opinion merging to the themes of the ICANN conference, particularly in the context of AI and machine learning:

- *Medical diagnosis:* specialists like radiologists and oncologists often provide critical opinions for treatment decisions. The proposed framework can effectively merge this knowledge, considering each expert's trustworthiness and certainty. This leads to more accurate and reliable treatment plans, especially in complex cases with few data.
- *Risk assessment:* Financial analysts frequently depend on expert opinions to evaluate investment risks. The framework can merge these insights while incorporating trust scores based on past performance and domain expertise. This refined approach enhances risk assessment models, resulting in better investment decisions and improved financial outcomes.
- *Product development and marketing:* companies can merge expert opinions about product features and market trends with consumer insights to select product development and marketing strategies. Integrating consumer opinions with expert insights ensures that products meet market needs while aligning with consumer preferences.

The paper layout is as follows. Section 1 briefly outlines the contemporary state of the art in the field of opinion merging. Section 2 introduces an example of the target problem and outlines the essential theory behind opinion merging. Section 3 adapts this theory to the specific context of merging opinions about mobile phone brands, providing a tailored framework for our analysis. Section 4 discusses modeller's trust in experts' opinions and their own level of certainty, highlighting the importance of these factors in the merging process. Section 5 presents the experimental results, comparison with weighted linear pooling [1] and discusses advantages of proposed method. Conclusion summarises the key findings from these experiments and suggests potential avenues for future research.

Conventions and notation

$\mathbb{N}_0$ set of natural numbers including 0 is denoted; $\theta \in \boldsymbol{\Theta}$ is a model parameter and $\boldsymbol{\Theta}$ is the set of parameters; $X = (\mathcal{B}, F_1, \ldots, F_n)$ is a discrete multivariate random variable, where $\mathcal{B}$ represents the brand ($b = 1, \ldots, m$) and each F_j ($j = 1, \ldots, n$) represents a feature score, with scores $f_{j,b}$ taking values in $\{1, \ldots, k\}$; $P, \mathcal{A}, \mathcal{S}$ represent probability distributions; The Kullback-Leibler divergence between distributions P and Q is expressed as $\mathcal{D}(P||Q)$; The subscript I indicates that a function is related to an agent, referred to as the modeller, while the subscript E signifies a relation to an external source of information, such as an expert; Additionally, the superscript $\cdot^U$ denotes the updated version of the corresponding variable. A strict preference relation between two outcomes is written as $a \succ b$;

2 Opinion Merging Task

In today's market, consumers face the challenge of choosing the best product while balancing their personal preferences with expert evaluations. Our goal is to support informed decision-making that reflects both individual beliefs and external insights/knowledge/experience. In this contribution, we focus on the evaluation of mobile phone brands by integrating personal preferences with expert assessments.

Imagine you're in the market for a new phone. You have your own ideas about which brand might be the best based on your experiences and what you've heard. However, you also have a tech-savvy friend who shares their opinion on which phone they believe is superior. Your aim is to combine your own thoughts with your friend's advice to make a more informed choice.

In this scenario, you act as the modeller, holding your own initial preferences about phone brands, while your friend's opinion serves as the external opinion. The process of merging these two perspectives can be seen as blending personal preferences with expert advice, ultimately leading to a decision that considers both viewpoints. You might weigh your friend's advice differently based on how much you trust their expertise. This approach ensures that your final decision reflects both your personal preferences and your trust in external insights.

We examine a scenario in which two independent agents each has an opinion about a random variable, denoted by X. One of these agents, the *modeller*, seeks to update their opinion based on two sources: data they have observed directly and the external opinion provided by the other agent, referred to as the *expert*. The objective is to understand how the modeller's updated opinion is formed by integrating these two sources of information.

The modeller's preference is represented as a joint distribution $P_I(X, \theta)$, where $\theta \in \Theta$ is a parameter representing the modeller's uncertain opinion. The distribution can be decomposed using the chain rule as:

$$P_I(X, \theta) = P_I(X|\theta) A_I(\theta), \tag{1}$$

where $P_I(X|\theta)$ is the modeller's opinion about X, and $A_I(\theta)$ is the modeller's prior belief about θ. The result of merging the modeller's opinion with the external opinion is represented by the joint distribution $P_I(X, \theta, A|P_E)$, which can be decomposed as follows:

$$P_I(X, \theta, A|P_E) = P_I(X|\theta, A, P_E) P_I(\theta|A, P_E) P_I(A|P_E).$$

Incorporating expert opinion $P_E(X)$ into $P_I(X, \theta)$ implies updating the prior distribution $A_I(\theta)$ to $A_I^U(\theta|P_E)$. Paper [8] address this problem by using Fully probabilistic design of hierarchical Bayesian models. It constructs a so-called hyper-prior $S_I(A|P_E)$ on unknown A and provides a justified way how to find its optimal $A_I(\theta|P_E)$. Then the optimal updated opinion derived using [8] reads:

$$P^U(X,\theta,A|P_E) = P_I(X)A_I^U(\theta|P_E)S_I^U(A|P_E), \text{ where}$$
$$S_I^U(A|P_E) \propto S_I(A|P_E)\exp\left(-\mathcal{D}(A||A_I^U)\right) \quad (2)$$

describes how to select appropriate update of the modeller prior (1) based on the expert's opinion. The updated belief of the modeller about Θ then reads

$$A_I^U(\theta|P_E) \propto P_I(\theta)\exp\left(\sum_x \ln(P_I(x|\theta))P_E(x)\right) \quad (3)$$

Meaning the resulting updated belief is a modification of the original prior $P_I(\theta)$ by exponential mixture of the modeller's original opinion and the opinion of an expert. The exponential function here appears from KLD minimization of opinion $P_I(X,\theta,A)$ (original opinion) and $P_I(X,\theta,A|P_E)$ (opinion modified by expert). Exact derivation can be found in [8].
The result of merging the modeller's opinion about X with the expert's $P_E(X)$ can be found as:

$$P_I^U(X) = \int_{\mathbf{\Theta}\times\mathcal{A}} P_I(X|\theta)A_I^U(\theta|P_E)S_I^U(A|P_E)d\theta dA.$$

3 Merging Opinions: Selecting a Mobile Phone Brand

To address the task of selecting mobile phone brands described in Sect. 2 we introduce a discrete multivariate random variable, $X = (\mathcal{B}, F_1, \ldots, F_n)$, where $\mathcal{B}$ represents the mobile phone brand and $F_j, j = \{1,\ldots,n\}$ denotes feature scores related to quality aspects such as camera performance, battery life, or display quality. The brands are indexed as b (ranging from 1 to m, with m being the total number of brands), and each feature score for brand b is represented as $f_{j,b}$ (with scores ranging from 1 to k, where k indicates the possible score range).

The modeller's opinion $P_I(X|\theta)$ reflects their belief about brands and corresponding features.

$$P_I(\mathcal{B}, F_1, \ldots, F_n|\theta) = \underbrace{P_I(\mathcal{B}|\theta)}_{\theta_\mathcal{B}}\underbrace{P_I(F_1|\mathcal{B},\theta)}_{\theta_{F_1}}\cdots\underbrace{P_I(F_n|\mathcal{B},\theta)}_{\theta_{F_n}}, \quad (4)$$

where $P_I(\mathcal{B}|\theta)$ indicates brand preference and $P_I(F_j|\mathcal{B},\theta)$ reflects opinion on specific features. Note the opinions are assumed conditionally independent. Let us consider $\theta = (\theta_\mathcal{B}, \theta_{F_1}, \ldots, \theta_{F_n})$ be an unknown parameter. Then prior belief $P_I(\theta)$, which captures the modeller's initial uncertainty, can be modelled using a Dirichlet distribution:

$$P_I(\theta) \propto \prod_b \theta_b^{n_b-1} \prod_{\{f_{j,b}\}_{j=1}^n} \theta_{f_{j,b}}^{n_{f_{j,b}}-1}. \quad (5)$$

In (5) weights $n_\bullet$ express the modeller's preference for brand n_b and particular feature $f_{j,b}$. The expert provides opinion $P_{E_i}(X), i \in \{1, \ldots, l\}$ about X

$$P_{E_i}(X) = P_{E_i}(\mathcal{B}) \prod_{j=1}^{n} P_{E_i}(F_j|\mathcal{B}). \tag{6}$$

In practice, each expert's opinion is often specific, assigning a particular brand and its associated feature scores directly. Therefore, we assume that the distributions representing experts' opinions about phone brands and features are degenerate.

Model $S_I(A|P_{E_i})$ (2) remains consistent regardless of the external opinion, reflecting a stable approach to incorporating new insights:

$$S_I(A|P_{E_i}) \propto \prod_b \prod_{\{f_{j,b}\}_{j=1}^n} \theta_{f_{j,b}}^{s_{f_{j,b}}-1}. \tag{7}$$

In (7) parameters $s_{f_{j,b}} \in \langle 0, 2\rangle$ quantify the modeller's trust in the expert's opinion and confidence in their own prior beliefs, specification of quantification is presented in Sect. 4 below. The updated prior, $A_I^U(\theta|P_{E_i})$, after merging with an expert's opinion (6) can be derived as:

$$A_I^U(\theta|P_{E_i}) \propto \frac{1}{\mathcal{N}_{A_I^U(\theta|P_{E_i})}} \theta_b^{n_b-1} \prod_{\{f_{j,b}\}_{j=1}^n} \theta_{f_{j,b}}^{V_{f_{j,b}}-1},$$

where updated weights $V_{f_{j,b}}$ incorporate both prior weights and the expert's influence:

$$V_{f_{j,b}} = n_{f_{j,b}} + s_{f_{j,b}} + \delta(b, b_{E_i})\delta(f_{j,b}, f_{E_i,j,b_{E_i}}) - 1.$$

The detailed steps of this derivation can be found in Appendix A. The final updated opinion of the modeller is represented as:

$$P_I^U(X = x) = \frac{n_{b_I}}{\sum_b n_b} \prod_{j=1}^{n} \frac{V_{f_{j,b_I}}}{\sum_{f_{j,b_I}} V_{f_{j,b_I}}}.$$

This formula encapsulates how the modeller balances their own prior beliefs with the influence of expert opinions.

After merging opinions, the updated brand preference is given by the Bayes rule:

$$P_I^U(\mathcal{B}|F_1, \ldots, F_n) = \frac{P_I^U(F_1|\mathcal{B}) \cdots P_I^U(F_n|\mathcal{B}) P_I^U(\mathcal{B})}{\sum_b P_I^U(F_{1,\mathcal{B}=b}|\mathcal{B}=b) \cdots P_I^U(F_{n,\mathcal{B}=b}|\mathcal{B}=b) P_I^U(\mathcal{B}=b)}. \tag{8}$$

Distribution (8) depends on the specific realizations of feature scores $F_1, \ldots, F_n$. These feature scores need to be selected appropriately to extract the modeller's preference for a particular brand $\mathcal{B}$. We assume that the modeller selects the feature score $f_{j,b}$ with the highest probability $P_I^U(F_{j,b}|\mathcal{B})$, i.e.

$$f_{j,b}^{opt} = \arg\max_{f_{j,b}}(P_I^U(F_{j,\mathcal{B}=b} = f_{j,b}|\mathcal{B} = b)), \tag{9}$$

for all brands $b \in \{1,\ldots,m\}$.

4 Trust and Certainty

Trust in the opinion provided by the expert P_{E_i} is $t_{E_i} \in \langle 0,1\rangle$. The lower value of t_{E_i} indicates a low trust in the opinion of experts E_i. Higher indicates otherwise. The trust value t_{E_i} can encapsulate various aspects of the relationship between the modeller and the expert E_i, such as friendship, prior knowledge of the expert's performance, or other subjective criteria set by the modeller. [1] For the purpose of experimentation, trust values t_{E_i} are predefined to specific values to illustrate the potential integration of trust t_{E_i} into the opinion merging process. The certainty of the modeller in their opinion $P_I(X)$, is represented by $c_{I,j,b} \in \langle 0,1\rangle$. A lower value of $c_{I,j,b}$ indicates diminished certainty in their opinion $P_I(X)$, while a higher value signifies the opposite. Inclusion of both trust t_{E_i} and opinion certainty $c_{I,j,b}$ into s_{f_j} is carried out through a linear combination: $s_{f_j} = t_{E_i} + c_{I,j,b}$. The preference sub-setting process dictates how the modeller chooses f_j that form the modeller's preference on mobile brands during the opinion merging. The process of preference sub-setting is accomplished using the modeller's score preference defined as $r_{I,j} \in \{1,2,\ldots,k\}$, representing the preference for the feature score F_j. This can be intuitively described as the modeller's inclination toward a specific range of scores f_j for a feature j across all mobile phone brands $b \in \mathcal{B}$. The score preference $r_{I,j}$ restricts the range for selecting scores $f_{j,b}^{\star}$ in (9) to $f_{j,b}^{\star} \geq r_{I,j}$.

Example (mobile phone selection task)

Let preference score $r_{I,j} = 4$ for quality of mobile phone display, the preferred feature score are highlighted in green:

Feature scores	1	2	3	4	5	6

. Consequently, this influences selection of the probabilities $P_I^U(F_j|\mathcal{B})$ entering the modeller's prior (8). These probabilities may assume a low value if the modeller's initial opinion P_I combined with expert's opinion P_{E_i} is outside the selection range defined by the score preference $r_{I,j}$.

5 Experiments

To showcase the functionality of the proposed opinion merging framework for the task of selecting a mobile phone based on the quality of its features, the following experimental setup on simulated data is considered. For comparison with the proposed method in this work, weighted linear pooling [1] of opinions

[1] It is important to note that the process of determining trust value t_{E_i} is not elaborated on in this work. Extraction of trust value can be achieved for example by similarity (dissimilarity measures) such solution can be found in [6,7,10].

was also implemented. The weights for the opinions of individual experts are given by the trust values t_{E_i}, while the weight for the opinion of the modeller is set to the certainty parameter $c_{I,j,b}$. Both parameters are defined as described in Sect. 4. Because the results obtained by applying both methods are similar, in the subsequent subsections we present the results of the opinion merging method proposed in this work, and discuss the comparison at the end of this section.

The simulated data were set up based on data collected from human participants in a survey about quality of features of several mobile brands. Showcase of proposed solution on simulated data was deemed to provide more insight into correctness of the proposed solution. The Python Jupyter notebook containing the experimental code and data is available at [12].

Let there be a person (the modeller) who is interested in buying a mobile phone. He chooses to be advised by their friends, colleagues and information on the internet. The modeller has a prior preference about mobile brands $\{B_1, B_2, B_3\}$ each having 3 features $\{F_1, F_2, F_3\}$. External information is in form of a score for each feature, the score scale is from 1 to 6, where 1 means the worst and 6 means the best. The experts E_i scores are set so that the solution to the selection process can be easily verified. The expert score tables applied in all Experiments are in Appendix C and are configured to the following ordering: $B_1 \succ B_2 \succ B_3$. The modeller's initial weights $n_{f_{j,b}}$ are set to 1 for each score. For simplicity the preference score $r_{I,j} = 4$ for all experiments, meaning the modeller prefers score of 4 or above for all the features. and brand preference $P_I(\mathcal{B})$ is configured to uniform distribution to ensure that the preference $P_I(\mathcal{B}|F_{1,\mathcal{B}}, \ldots, F_{n,\mathcal{B}})$ is not influenced by the modeller's overall preference of brand. That is, only the preference based on features is considered.

5.1 Opinion Merging and Preference Subsetting

This section illustrates the influence of expert opinion on the modeller's brand preference. In the following experiment, the opinions are merged without considering modeller's trust in expert opinions and without considering certainty of the modeller into his opinion. These factors will be investigated in the subsequent experiments.

Experiment 1: The aim of this experiment is to test if this ordering is shifted towards B_1 in the final updated brand preference (8). The modeller's prior preference on brands is given by the ordering: $B_2 \succ B_1 \succ B_3$ and can be found in Table: 6. Experts slightly prefer B_1 to B_2, which is a setup for the later demonstration of the trust value t_{E_i} (see Sect. 5.2). The result of final preference of mobile brands is in Fig. 1. The results suggest that the experts influenced the modeller to favour the B_1 brand. Which is the expected outcome.

5.2 Inclusion of Trust

In this section, we examine the integration of trust t_{E_i} for each expert E_i's opinion. In the following two experiments, the opinions are merged without considering modeller's certainty in his opinion. For simplicity, the setup of modeller's initial scores for the following two experiments is the same as in the first experiment.

Experiment 2: We start by configuring low trust values t_{E_i} for experts reacting to the B_1 brand and high trust values t_{E_i} for experts reacting to the B_2 and B_3, specific trust scores can be found in Table 4. The expected result is that the preferred brand should be the B_2, since it has similar score values provided by experts E_i as B_1, with the B_2 being slightly less favoured. The result of final preference of mobile brands is in Fig. 2. The result of this experiment complies with the described expectation.

Experiment 3: Setting the trust t_{E_i} for experts reacting to the B_1 and B_2 brands to low values and to high values for B_3, specific trust scores can be found in Table 5. The result of final preference of mobile brands is in Fig. 3. Both experiments complied with the described expected results. The trust process works as expected.

5.3 Inclusion of Certainty

To better demonstrate the proposed solution, trust values t_{E_i} are intentionally excluded. When the trust value is not set, it is equivalent to setting the trust t_{E_i} to the maximum value of 1.

Experiment 4: This experiment illustrates the impact of the certainty $c_{I,j,b}$ on the final ordering of brands. The modeller's opinion is configured to prefer the brands in the following order: $B_2 \succ B_3 \succ B_1$ and is presented in Table 7. Setting the maximum certainty of opinion $c_{I,j,b} = 1$ in the low scores of the modeller for the B_1 brand should lead to the preference for the B_2 as the main brand, with B_3 being the second most preferred brand. The modeller's certainty $c_{I,j,b}$ for each respective brand B_1, B_2, B_3 is:

Certainty	100%	40%	10%

. This can be intuitively interpreted as the modeller being less receptive to the advice of experts E_i about B_1. The result of final preference of mobile brands is in Fig. 4. The result of this experiment complies with the described expectation.

Experiment 5: In this experiment the modeller's opinion is configured to prefer the brands in the following order: $B_3 \succ B_2 \succ B_1$ and can be found in Table 8. Maximum certainty $c_{I,j,b}$=1 is applied to the modeller's high scores for B_3 brand and should lead to the highest preference for the B_3 brand. The modeller's certainty $c_{I,j,b}$ for each respective brand B_1, B_2, B_3 is:

Certainty	20%	20%	100%

. High values of certainty $c_{I,j,b}$ indicate that the modeller is very confident in brand B_3 and, therefore, unlikely to change his opinion, even if the experts overwhelmingly prefer other brands. The result of final preference of mobile brands is in Fig. 5. The result of this experiment complies with the described expectation.

5.4 Comparison of Linear Pooling and FPD

The resulting brand preference $P_I(\mathcal{B} \mid F_{1,\mathcal{B}}, \ldots, F_{n,\mathcal{B}})$ after applying linear opinion pooling was essentially the same as that obtained using the proposed method. The simple average difference in the preference for individual brands across all experiments are almost negligible: | Difference | 0.041% | 0.023% | 0.034% | . This outcome is expected, as the implementation of both methods in this work is nearly identical. The difference arises from the initialization of opinions. In the FPD-based method, the modeller's opinion is initialized with non-zero weights using a Dirichlet prior, and the resulting probabilities are calculated as mean estimates of the parameters. In contrast, linear opinion pooling constructs the probability distribution directly, without parameter estimation. This result validates the proposed method. The main distinction is that the method presented here provides a theoretically sound Bayesian formulation, allowing for online updating of the modeller's opinion in a dynamic setting, which is subject of future research, see [4]. Achieving a similar effect with the linear pooling method would require recalculating the opinion mixture each time a new opinion needs to be incorporated to ensure theoretical correctness.

6 Conclusion

This work addresses the challenge of forming a new opinion by integrating insights from experts, the agent's prior beliefs, their trust in those experts, and their confidence in their own opinions. By adopting a *single-agent perspective* and modelling opinions probabilistically, we framed the opinion merging task as a *decision-making* problem. To validate the proposed approach, we employed the selection of a mobile phone brand with desired features as a practical test case and performed weighted opinion pooling [1] of the same opinions for baseline comparison. The resulting solution is represented as a posterior distribution of mobile brands, which captures the merged opinions of both the modeller and the experts. This distribution respects the agent's prior beliefs, their preferences for specific features, and the importance of those features. Data were collected from human participants. Given the lack of a definitive ground truth for merged opinions, we designed a series of simulated experiments based on real-world data that can be assessed through logic and intuition. The results obtained are robust and self-explanatory. The results from the proposed method are almost identical to results from weighted linear opinion pooling, which validates the method. Main advantage of the proposed solution is theoretically sound Bayesian formulation which allows on-line updating of opinions in dynamic setting. Several key questions remain for future research:

- How can we effectively learn an agent's trust in an expert's opinion based on previously observed data, external knowledge, and varying preferences?
- How can we model the dynamics of trust, particularly its evolution with growing experience and changing agent's preferences?

– In what ways does confidence in the provided information shape the formation of opinions?

Our long-term vision is to combine the learning of preferences with opinion merging, ultimately creating a reliable, analytically transparent human-in-the-loop support system for interactions within social or customer networks. This approach, in combination with decision-making systems, has the potential to enhance and transparetize decision-making processes and foster more effective communication across diverse contexts and domains such as merging of opinions of medical specialist, company managers, policy-makers or autonomous systems.

Acknowledgments. The results and knowledge included herein have been obtained owing to support from the EU-funded IAMI project, Grant agreement ID: 101168272 and the institutional grant no. 2025A1013, Internal Grant Agency of the Faculty of Economics and Management, Czech University of Life Sciences Prague.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

Appendix A

Let us substitute $P_I(X)$ (4) $P_I(\theta)$ (5) and $P_{E_i}(X)$ (6) into $A_I^U(\theta|P_{E_i})$ (3). Substituting $P_I(\theta)$ (5) and simple equivalent transformations gives:

$$A_I^U(\theta|P_{E_i}) = \prod_b \prod_{\{f_{j,b}\}_{j=1}^n} \theta_b^{n_b-1}\theta_{f_{j,b}}^{n_{f_{j,b}}-1} \exp\left[\sum_{x\in X} \ln(P_I(x|\theta))P_{E_i}(x)\right]. \tag{10}$$

Substituting (4) and (6) into (10), where expert provides, among other, $P_{E_i}(\mathcal{B})$. This distribution does not represent the expert's preference for brand $\mathcal{B}$. Therefore, can be omitted during the substitution of (6).

$$\begin{aligned}A_I^U(\theta|P_{E_i}) \propto \prod_b \prod_{\{f_{j,b}\}_{j=1}^n} \theta_b^{n_b-1}\theta_{f_{j,b}}^{n_{f_{j,b}}-1} \exp\Bigg[\sum_{x\in X} \ln(\theta_b\theta_{f_{1,b}}\cdots\theta_{f_{n,b}}) \\ \times P_{E_i}(F_{1,\mathcal{B}=b} = f_{1,b}|\mathcal{B}=b)\cdots P_{E_i}(F_{n,\mathcal{B}=b} = f_{n,b}|\mathcal{B}=b)\Bigg].\end{aligned} \tag{11}$$

As noted, external opinion $P_{E_i}(X)$ degenerates, then (11) reads

$$A_I^U(\theta|P_{E_i}) \propto \prod_b \prod_{\{f_{j,b}\}_{j=1}^n} \theta_b^{n_b-1}\theta_{f_{j,b}}^{n_{f_{j,b}}-1} \theta_{f_{E_i,1,b_{E_i}}}\cdots\theta_{f_{E_i,n,b_{E_i}}}. \tag{12}$$

Next, we will make the following adjustment in (12):

$$A_I^U(\theta|P_{E_i}) \propto \prod_b \prod_{\{f_{j,b}\}_{j=1}^n} \theta_b^{n_b-1} \theta_{f_{j,b}}^{n_{f_{j,b}}+\delta(b,b_{E_i})\delta(f_{j,b},f_{E_i,j,b_{E_i}})-1}.$$

Next, we multiply by the modeller's strategy $S_I(A|P_{E_i})$ specified in (7), since it serves as a hyper-prior for the distribution $A_I^U(\theta|P_{E_i})$ and move the relevant exponential terms under one exponent:

$$A_I^U(\theta|P_{E_i}) \propto \prod_b \prod_{\{f_{j,b}\}_{j=1}^n} \theta_b^{n_b-1} \theta_{f_{j,b}}^{n_{f_{j,b}}+s_{f_{j,b}}-1+\delta(b,b_{E_i})\delta(f_{j,b},f_{E_i,j,b_{E_i}})-1} \tag{13}$$

We introduce the following definition of the updated weights.

$$V_{f_{j,b}} = n_{f_{j,b}} + s_{f_{j,b}} + \delta(b, b_{E_i})\delta(f_{j,b}, f_{E_i,j,b_{E_i}}) - 1,$$

where b is a brand, b_{E_i} denotes brand selected by i-th expert, $f_{j,b}$ is score of j-th feature $f_{E_i,j,b_{E_i}}$ is a score selected by i-th expert for brand b_{E_i}. Then (13) reads

$$A_I^U(\theta|P_{E_i}) \propto \prod_b \theta_b^{n_b-1} \prod_{\{f_{j,b}\}_{j=1}^n} \theta_{f_{j,b}}^{V_{f_{j,b}}-1}, \tag{14}$$

To obtain the final form of updated prior (3), we need to calculate the normalisation constant. This is done by noticing that the final distribution is a product of the Dirichlet distributions. Therefore, the normalisation constant is a product of normalisation constants.

$$\mathcal{N}_{A_I^U(\theta|P_{E_i})} = \prod_b \frac{\Gamma(n_b)}{\Gamma(\sum_b n_b)} \prod_j \frac{\prod_{f_{j,b}} \Gamma(V_{f_{j,b}})}{\Gamma(\sum_{f_{j,b}} V_{f_{j,b}})}, \tag{15}$$

where $j \in \{1, \ldots, n\}$, $b \in \{1, \ldots, m\}$ and Γ is the Gamma function.

After merging with the expert's opinion $P_{E_i}(X)$ (6), the final form of the updated prior $A_I^U(\theta|P_{E_i})$ is :

$$A_I^U(\theta|P_{E_i}) = \frac{1}{\mathcal{N}_{A_I^U(\theta|P_{E_i})}} \prod_b \prod_{\{f_{j,b}\}_{j=1}^n} \theta_b^{n_b-1} \theta_{f_{j,b}}^{V_{f_{j,b}}-1}. \tag{16}$$

Appendix B

The modeller's updated opinion can be written as:

$$P_I^U(X) = \int_{\Theta} P_I^U(X,\theta)d\theta = \int_{\Theta} P_I(X|\theta) A_I^U(\theta|P_{E_i})d\theta. \tag{17}$$

In the second equality, $A_I^U(\theta|P_{E_i})$ is the updated prior specified in (16). Next, we substitute the modeller's opinion (4) and updated prior $A_I^U(\theta|P_{E_i})$ (16) into

(17). Substituting $P_I(X|\theta)$ (4) into (17), using assumption of conditional independence as in (4) and equivalent transformations gives:

$$P_I^U(X=x) = \frac{1}{\mathcal{N}_{A_I^U(\theta|P_{E_i})}} \int_{\Theta} \prod_b \theta_b^{n_b+\delta(b,b_I)-1} \prod_{\{f_{j,b}\}_{j=1}^n} \theta_{f_{j,b},b}^{V_{f_{j,b}}+\delta(b,b_I)\delta(f_{j,b},f_{j,b_I})-1} d\theta. \tag{18}$$

Similarly to the calculation in (14), the integral is a product of Dirichlet distributions

$$P_I^U(X) = \frac{n_{b_I}}{\sum_b n_b} \prod_j \frac{V_{f_{j,b_I}}}{\sum_{f_{j,b_I}} V_{f_{j,b_I}}}. \tag{19}$$

Appendix C

Simulated expert opinion tables for each brand (Tables 1, 2 and 3).

Table 1. Brand 1.

Experts	Feature 1	Feature 2	Feature 3
1	4	3	4
2	5	3	5
3	5	6	5
4	6	5	3
5	6	6	6
6	5	6	5
7	6	6	5
8	6	3	4
9	4	5	4
10	6	4	3

Table 2. Brand 2.

Experts	Feature 1	Feature 2	Feature 3
1	5	5	5
2	5	6	5
3	3	4	4
4	3	4	5
5	4	5	5
6	5	6	4
7	6	6	6
8	5	6	6
9	4	3	4
10	4	6	3

Table 3. Brand 3.

Experts	Feature 1	Feature 2	Feature 3
1	3	4	3
2	3	3	4
3	4	3	4
4	3	3	3
5	5	4	3
6	3	5	5
7	4	5	6
8	4	3	2
9	3	4	3
10	4	3	4

Trust levels assigned to experts by the modeller for individual experiments.

Table 4. Experiment 2.

Expert	Brand 1	Brand 2	Brand 3
1	0.8	0.9	0.9
2	0.9	0.9	0.7
3	0.3	0.3	0.8
4	0.7	0.3	0.8
5	0.2	0.7	0.9
6	0.2	0.7	0.8
7	0.3	0.9	0.7
8	0.5	0.9	0.7
9	0.4	0.4	0.8
10	0.9	0.4	0.9

Table 5. Experiment 3.

Expert	Brand 1	Brand 2	Brand 3
1	0.8	0.3	0.5
2	0.9	0.3	0.7
3	0.5	0.9	0.9
4	0.9	0.9	0.2
5	0.2	0.5	0.9
6	0.2	0.5	0.8
7	0.2	0.2	0.9
8	0.5	0.2	0.7
9	0.4	0.8	0.8
10	0.9	0.8	0.9

The modeller's opinion setup for individual experiments.

Table 6. Experiments 1–3.

	Feature 1	Feature 2	Feature 3
Brand 1	5	4	5
Brand 2	6	5	6
Brand 3	4	4	3

Table 7. Experiment 4.

	Feature 1	Feature 2	Feature 3
Brand 1	3	4	1
Brand 2	6	5	6
Brand 3	4	4	3

Table 8. Experiment 5.

	Feature 1	Feature 2	Feature 3
Brand 1	3	1	1
Brand 2	2	3	1
Brand 3	6	5	5

Updated brand preferences of the modeller across five experiments.

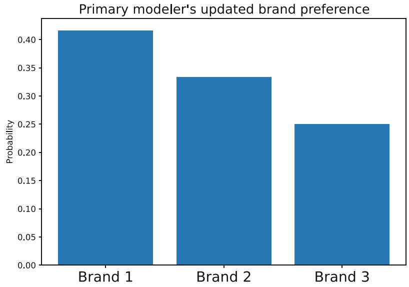

Fig. 1. Experiment 1.

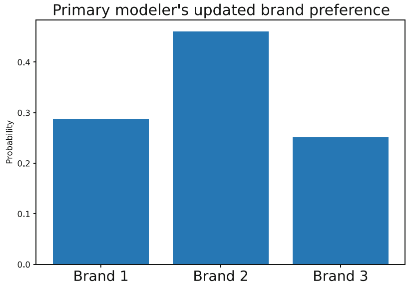

Fig. 2. Experiment 2.

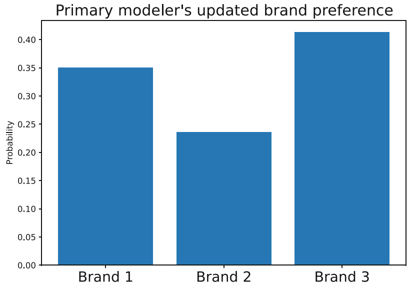

Fig. 3. Experiment 3.

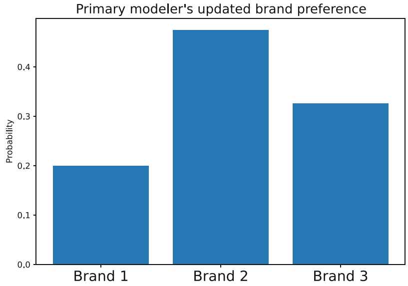

Fig. 4. Experiment 4.

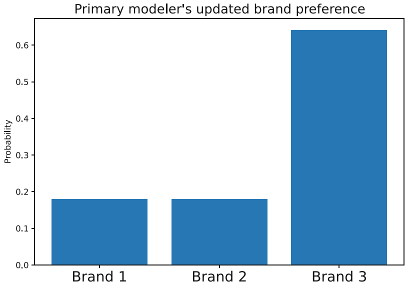

Fig. 5. Experiment 5.

References

1. DeGroot, M.H., et al.: Optimal linear opinion pools. Manag. Sci. (1991)
2. Garg, A., et al.: Generalized opinion pooling. In: AI&M (2004)
3. Genest, C., et al.: Allocating the weights in the linear opinion pool. J. Forecast. (1990)
4. Karny, M.: Optimized geometric pooling of probabilities for information fusion and forgetting. Automatica (2025)
5. Liu, P., et al.: Social trust-driven consensus reaching model for multi-attribute group decision making: exploring social trust network completeness. TFS (2023)
6. O'Donovan, J., et al.: Trust in recommender systems. In: Proceedings of the 10th International Conference on Intelligent User Interfaces (2005)
7. O'Doherty, D., et al.: Trust-based recommendation: an empirical analysis. In: Proceedings of the 6th ACM Workshop on Social Network Mining and Analysis (2012)
8. Quinn, A., et al.: Optimal design of priors constrained by external predictors. Int. J. Approximate Reasoning 153 (2017)
9. Romeijn, J.W.: An interpretation of weights in linear opinion pooling. Episteme (2020)
10. Roy, F., et al.: Comparative analysis of different trust metrics of user-user trust-based recommendation system. Comput. Sci. (2022)
11. Rufo, M.J., et al.: Merging experts' opinions: a bayesian hierarchical model with mixture of prior distributions. Eur. J. Oper. Res. (2010)
12. Ružejnikov, J.: Targeted trust based merging of customers opinions experiments: v1.0 (2025). https://doi.org/10.5281/zenodo.15738051
13. Sarwate, A.D.: Merging opinions by social sampling of posteriors. In: Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton) (2012)
14. Ya, W., et al.: Consensus model of social network group decision-making based on trust relationship among experts and expert reliability. J. Syst. Eng. Electron. (2023)