Agent Teams for Design Problems

Leandro Soriano Marcolino[1], Haifeng Xu[1], David Gerber[3], Boian Kolev[2], Samori Price[2], Evangelos Pantazis[3], and Milind Tambe[1]

[1] Computer Science Department, University of Southern California, CA, USA
[2] Computer Science Department, California State University, Dominguez Hills, USA
[3] School of Architecture, University of Southern California, CA, USA
Department of Civil Engineering, University of Southern California, CA, USA
{sorianom,haifengx,dgerber}@usc.edu,
{bkolev1,sprice25}@toromail.csudh.edu, {epantazi,tambe}@usc.edu

Abstract. Design imposes a novel social choice problem: using a team of voting agents, maximize the number of optimal solutions; allowing a user to then take an aesthetical choice. In an open system of design agents, team formation is fundamental. We present the first model of agent teams for design. For maximum applicability, we envision agents that are queried for a single opinion, and multiple solutions are obtained by multiple iterations. We show that diverse teams composed of agents with different preferences maximize the number of optimal solutions, while uniform teams composed of multiple copies of the best agent are in general suboptimal. Our experiments study the model in bounded time; and we also study a real system, where agents vote to design buildings.

Keywords: Collaboration, Distributed AI, Team Formation

1 Introduction

Teams of voting agents are a power tool for finding the optimal solution in many applications [15, 1, 16, 18, 10], as there are theoretical guarantees in finding one optimal choice [5]. For design problems, however, finding one optimal solution is not enough, and we actually want to find as many optimal solutions as possible, allowing a human to choose according to her aesthetical taste. Even if a user does not want to consider too many solutions, they can be filtered and clustered [7], allowing her to easily make an aesthetical choice. Hence, a system of voting agents that produces a unique optimal solution is insufficient; and we propose the novel social choice problem of maximizing the number of optimal alternatives found by a voting system. As ranked voting may suffer from noisy rankings when using existing agents [11], we study multiple voting iterations.

Traditionally, social choice studies the optimality of voting rules, assuming certain noise models for the agents, and rankings composed of a linear order over alternatives [5, 4, 14]. Hence, there is a single optimal choice, and a system is successful if it can return that optimal choice with high probability. More recently, several works have been considering cases where there is a partial order over alternatives [22, 19], or when the agents output pairwise comparisons instead of rankings [6]. However, these works still focus on finding an optimal alternative, or a fixed-sized set of optimal alternatives (where the size is known beforehand). Therefore, they still provide no help in finding the maximum set of optimal solutions. Moreover, they assume agents that are able to output comparisons among all actions with fairly good precision, and the use of multiple voting iterations has never been studied. When considering agents with different preferences, the field is focused on verifying if voting rules satisfy a set of axioms that are considered to be important to achieve fairness [17].

In this work we offer a completely different perspective: we show that, unless we have an idealized agent, we only maximize the number of optimal solutions if we have agents with different preferences. Motivated by the need of selecting agents from an open system, for greater applicability we only consider agents that output a single action. We present a theoretical study of which teams are desirable for design problems, and how their size may effect optimality. We show that, contrary to traditional social choice models, increasing the team size may significantly harm performance; and that a diverse team of agents with different preferences is highly desirable for achieving optimality. In doing so, we draw a novel connection between social choice and *number theory*; allowing us to show, for example, that the optimal diverse team size is constant with high probability, and a prime number of optimal actions may impose problems. We present synthetic experiments to further study our model, providing realistic insights into what happens with bounded computational time. Finally, we present experiments in a highly relevant domain: architecture design, where we show teams of agents that vote to design energy-efficient buildings. Hence, this is the first work exploring and showing the potential of voting systems in being *creative*, by actually creating new alternatives from the opinions of existing agents.

2 Related Work

As mentioned, traditional works in social choice concern finding a correct ranking in domains where there is a unique optimal decision [5, 4, 14, 20]. Recent works, however, are considering more complex domains. Xia and Conitzer (2001) [22] study the problem of finding k optimal solutions, where k is known beforehand, by aggregating rankings from each agent. However, not only do they need strong assumptions about the quality of the rankings of such agents, but they also show that calculating the MLE from the rankings is an NP-hard problem.

Procaccia *et al.* [19] study a similar perspective, where the objective is to find the top k options given rankings from each agent, where, again, k is known in advance. However, in their case, they assume there still exists one unique truly optimal choice, hidden among these top k alternatives. Elkind and Shah (2014) [6] study the case where instead of rankings, the voters output pairwise comparisons among all actions, which may not follow transitivity. However, their final objective is still to pick a single winner.

Finally, outputting a full comparison among all actions can be a burden for an agent [3]. Jiang *et al.* (2014) [11] show that actual agents can have very noisy rankings, and therefore do not follow the assumptions of previous works in social choice. Hence, as any agent is able to output at least one action (i.e., a single vote), we study here systems where agents vote across multiple iterations.

3 Design Domains

We consider in this work domains where the objective is to find the highest number of optimal solutions. We show that design is one of such domains. One of the most common computational design approaches is to use *parametric designs* [21, 9, 7], where a human designer creates an initial design of a product using *computer-aided design* tools. However, instead of manually deciding all aspects of the product, she leaves *free parameters*, whose values can be modified to change the design. This approach is used because the number of different possibilities that a human can manually create while looking for optimality is limited, so a computer system is used to refine the design and find optimal solutions.

Design problems are in general multi-objective [12], since a product normally must be optimized across different factors. For example, a product should have a low cost, but at the same time high quality, two highly-contradictory objectives. Hence, there are a large number of optimal solutions, all tied in a *pareto frontier*. For the computational system, these optimal solutions are all equivalent. However, a human may dynamically decide to value some factor over another, and/or pick the option that most pleases her own aesthetical taste or the one of the target public/client.

Note that choosing a design according to aesthetics is an undefined problem, since there are no formal definitions to compare among different options. Hence, the best that a system can do is to provide a human with a large number of optimal solutions (according to other measurable factors), allowing her to freely decide among equally optimal solutions — but most probably with not equal aesthetical qualities.

Therefore, it is natural that in design problems we are going to have many possible solutions, and we want to find as many optimal ones as possible. In fact, there are many benefits in discovering a large number of optimal solutions:

Knowledge Does not Hurt: Having more optimal solutions to choose from is never worse than having less. For example, if a designer has enough time to analyze only x solutions, she can do so with a system that provides more than x optimal solutions by sampling the exact amount that she desires. However, she will never be able to do so with a system that provides less than x optimal solutions. Moreover, we can easily identify and eliminate solutions that are similar by applying clustering and analysis techniques [7], so that every solution that the human looks at is meaningful.

Knowledge Increases Confidence in Optimality: In general design problems, the true pareto frontier is unknown. Genetic Algorithms are widely used in order to *estimate* it. The only knowledge available for the system to evaluate the optimality is in comparison with the other solutions that are also being evaluated during the optimization process [13]. Many apparently "optimal" solutions are actually discovered to be sub-optimal as we find more solutions. Hence, finding a higher number of optimal solutions decreases the risk of a designer picking a wrong choice that was initially outputted as "optimal".

Knowledge Increases Aesthetical Qualities: If a human has a larger set of optimal solutions to choose from, there is a greater likelihood that at least one of these solutions is going to be of high aesthetical quality according to her preferences, or the ones of the target public.

Knowledge Increases Diversity of Options: In general, when a system x has more optimal solutions available than a system y, it does not necessarily imply that the solutions in the system x are more similar, while the optimal solutions in y are more different/diverse. In fact, all things equal, the greater the amount of optimal solutions, the higher the likelihood that we have more diverse solutions available.

4 Agent Teams

We present our theory of agent teams for design problems. Consider a team that vote together at each possible decision point of the design of a product (for example, they may vote for the value of each parameter, in a parametric design). Hence, let $\mathbf{\Phi}$ be a set of agents ϕ, and $\mathbf{\Omega}$ a set of world states ω. Each ω has an associated set of possible actions $\mathbf{A}_\omega$. At each world state, each agent ϕ_i outputs an action a_j, an optimal action according to the agent's imperfect evaluation – which may or may not be a true optimal action. Hence, there is a probability p_j that the agent outputs a certain action a_j. The teams take the action decided by *plurality voting* (ties are broken randomly). We assume that the world states are independent, and by taking an optimal action at all world states we find an optimal solution for the entire problem.

In this paper our objective goes beyond finding one optimal solution, we want to maximize the number of optimal solutions that we can find. For greater applicability, we consider here agents that output a single action. Hence, we generate multiple solutions by re-applying the voting procedure across all world states multiple times (which are called *voting iterations* – one *iteration* goes across all world states, forming one solution). Formally, let $\mathbf{S}$ be the set of (unique) optimal solutions that we find by re-applying the voting procedure through z iterations. Our objective is to maximize $|\mathbf{S}|$. We will show that, under some conditions, we can achieve that when $z \to \infty$ (we study bounded time in Section 5).

We consider that at each world state ω there is a subset $\mathbf{Good}_\omega \subset \mathbf{A}_\omega$ of optimal actions in ω. An optimal solution is going to be composed by assigning any $a_j \in \mathbf{Good}_\omega$ in world state ω – for all world states. Conversely, we consider the complementary subset $\mathbf{Bad}_\omega \subset \mathbf{A}_\omega$, such that $\mathbf{Good}_\omega \cup \mathbf{Bad}_\omega = \mathbf{A}_\omega, \mathbf{Good}_\omega \cap \mathbf{Bad}_\omega = \emptyset$. We drop the subscripts ω when it is clear that we are referring to a certain world state.

One fundamental problem is selecting which agents should form a team. By the classical voting theories, one would expect the best teams to be uniform teams composed of multiple copies of the best agent [5, 14]. Here we show, however, that for design problems uniform teams need very strong assumptions to be optimal, and in most cases they actually converge to always outputting a single solution – an undesirable outcome. However, diverse teams are optimal as long as the team's size grows *carefully*, as we explain below in Theorem 1.

We call a team *optimal* when $|\mathbf{S}| \rightarrow \prod_{\omega} |\mathbf{Good}_{\omega}|$ as $z \rightarrow \infty$, and all optimal actions are chosen by the team with the same probability. Otherwise, even though the team still produces all optimal solutions, it would tend to repeat already generated solutions whose probability is higher. Since in practice there are time bounds, such condition is fundamental to have as many solutions as possible in limited time.

We first consider agents that are independent and identically distributed. Let p_j^{Good} be the probability of voting for $a_j \in \mathbf{Good}$, and p_k^{Bad} be the probability of voting for $a_k \in \mathbf{Bad}$. Let $n = |\boldsymbol{\Phi}|$ be the size of the team, and N_l be the number of agents that vote for a_l in a certain voting iteration. If $\forall a_j \in \mathbf{Good}, a_k \in \mathbf{Bad}$, $p_j^{Good} > p_k^{Bad}$, the team is going to find one optimal solution with probability 1 as $n \rightarrow \infty$, as we show in the following observation:

Observation 1 *The probability of a team outputting one optimal solution goes to* 1 *as* $n \rightarrow \infty$, *if* $p_j^{Good} > p_k^{Bad}$, $\forall a_j \in \mathbf{Good}, a_k \in \mathbf{Bad}$.

Proof. Note that as the agents are independent and identically distributed, we can model the process of pooling the opinions of n agents as a multinomial distribution with n trials (and the probability of any class k of the multinomial corresponds to the probability p_k of voting for an action a_k).

Hence, for each action a_l, the expected number of votes is given by $E[N_l] = n \times p_l$. Therefore, by the *law of large numbers*, if $p_j^{Good} > p_k^{Bad}$ $\forall a_j \in \mathbf{Good}, a_k \in \mathbf{Bad}$, we have that $N_j > N_k$. Hence, the team will pick an action $a_j \in \mathbf{Good}$, in all world states, if n is large enough (i.e., $n \rightarrow \infty$). ■

However, with a team made of copies of the same agent, the system is likely to lose the ability to generate new solutions as n increases. If, for each ω, we have an action a_m^{ω} such that $p_m^{Good} > p_j^{Good}$ $\forall a_m^{\omega} \neq a_j^{\omega}$, the team converges to picking only action a_m^{ω} (Proposition 1 below). Hence, $|\mathbf{S}| = 1$, which is a very negative result. Therefore, contrary to traditional social choice, here it is not the case that increasing the team size always improves performance.

Let $p_{Good} = \sum_j p_j^{Good}$ be the probability of picking any action in **Good**. We re-write the probability of an action a_j^{Good} as: $p_j^{Good} = \frac{p_{Good}}{|\mathbf{Good}|} + \lambda_j$, where $\sum_j \lambda_j = 0$. Let $\boldsymbol{\lambda}^{+}$ be the set of $\lambda_j > 0$. Let λ^{High} be the maximum possible value for $\lambda_j \in \boldsymbol{\lambda}^{+}$, such that the relation $p_j^{Good} > p_k^{Bad}$, $\forall a_j \in \mathbf{Good}, a_k \in \mathbf{Bad}$ is preserved. We show that when $z \rightarrow \infty$, $|\mathbf{S}|$ is the highest as $\max \boldsymbol{\lambda}^{+} \rightarrow 0$, and the lowest (i.e., one) as $\min \boldsymbol{\lambda}^{+} \rightarrow \lambda^{High}$.

Proposition 1. *The maximum value for* $|\mathbf{S}|$ *is* $\prod_{\omega} |\mathbf{Good}_{\omega}|$. *When* $z, n \rightarrow \infty$, *as* $\max \boldsymbol{\lambda}^{+} \rightarrow 0$, $|\mathbf{S}| \rightarrow \prod_{\omega} |\mathbf{Good}_{\omega}|$. *Conversely, as* $\min \boldsymbol{\lambda}^{+} \rightarrow \lambda^{High}$, $|\mathbf{S}| \rightarrow 1$.

Proof. As $\max \boldsymbol{\lambda}^{+} \to 0$, $\lambda_j \to 0$, $\forall a_j$. Hence, $E[N_j] \to n \times \frac{p_{Good}}{|\mathbf{Good}|}$, $\forall a_j \in \mathbf{Good}$. Because ties are broken randomly, at each world state ω, each $a_j \in \mathbf{Good}_\omega$ is selected by the team with equal probability $\frac{1}{|\mathbf{Good}_\omega|}$. As $E[N_j] = E[N_k]\ \forall a_j, a_k \in \mathbf{Good}$, we have that at each ω it is possible to choose $|\mathbf{Good}_\omega|$ different actions. Hence, there are $\prod_\omega |\mathbf{Good}_\omega|$ possible combinations of solutions. At each voting iteration, ties are broken at each ω randomly, and one possible combination is generated. As $z \to \infty$, eventually we cover all possible combinations, and $|\mathbf{S}| \to \prod_\omega |\mathbf{Good}_\omega|$.

Conversely, as $\min \boldsymbol{\lambda}^{+} \to \lambda^{High}$, $E[N_j] \to n \times p_j^{Good}$ for one fixed a_j such that $p_j^{Good} > p_k^{Good}, \forall a_j \neq a_k \in \mathbf{Good}$ (and, consequently, $E[N_j] > E[N_k]$), at each ω. Hence, there is no tie in any world state, and the team picks a fixed a_j at each world state. Therefore, even if $z \to \infty$, $|\mathbf{S}| \to 1$. ■

Therefore, uniform teams need a very strong assumption to be optimal: the probability of voting for optimal actions must be uniformly distributed over all optimal actions ($\max \boldsymbol{\lambda}^{+} \to 0$). We show that, alternatively, we can use agents with different "preferences" (i.e., "diverse" agents), to maximize $|\mathbf{S}|$. We consider here agents that have about the same ability in problem-solving, but they prefer different optimal actions. As the agents have similar ability, in order to simplify the analysis we consider the probabilities to be the same across agents, except for the actions in **Good**, as each agent ϕ_i has a subset $\mathbf{Good^i} \subset \mathbf{Good}$ consisting of its preferred actions (which are more likely to be chosen than other actions). We denote by p_{ij} the probability of agent ϕ_i voting for action a_j. $\forall a_j \in \mathbf{Good}^i$, let $p_{Good^i} = \sum_j p_{ij}$, $p_{ij} = \frac{p_{Good^i}}{|\mathbf{Good}^i|}$, and $p_{ij} > p_{ik}, \forall a_k \notin \mathbf{Good^i}$. $\mathbf{Good}^i \cap \mathbf{Good}^l$ (of agents ϕ_i and ϕ_l) is not necessarily $\emptyset$. Consider we can draw diverse agents from a distribution $\mathcal{F}$. Each agent ϕ_i has $r < |\mathbf{Good}|$ actions in its $\mathbf{Good}^i$, and we assume that all actions in **Good** are equally likely to be selected to form $\mathbf{Good}^i$ (since they are all equally optimal). Note that each agent can even prefer a single action ($r = 1$), so this is a realistic assumption. We show that by drawing n agents from $\mathcal{F}$, the team is optimal for large n with probability 1, as long as n is a multiple of a divisor (> 1) of each $|\mathbf{Good}_\omega|$. We also show that the minimum necessary optimal team size is constant with high probability as the number of world states grows. We start with the following proposition:

Proposition 2. *If a team of size n is optimal at a world state, then* $gcd(n, |\mathbf{Good}|) > 1$ *(n and* $|\mathbf{Good}|$ *are not co-prime).*

Proof. Prove by contradiction. By the optimality requirements we must have $nr = k|\mathbf{Good}|$, where k is a constant $\in \mathbb{N}_{>0}$ representing the number of agents that have a given action a_j in its $\mathbf{Good}^i$ – note that it must be the same for all optimal actions. If n and $|\mathbf{Good}|$ are co-prime, then it must be the case that r is divisible by $|\mathbf{Good}|$. However, this yields $r \geq |\mathbf{Good}|$, which contradicts our assumption. Therefore, n and $|\mathbf{Good}|$ are not co-prime.■

This implies hard restrictions for world states where $|\mathbf{Good}|$ is prime, or for teams with prime size n: if n is prime, $|\mathbf{Good}|$ must be a multiple of n; and if $|\mathbf{Good}|$ is prime, n must be a multiple of $|\mathbf{Good}|$.

Now we consider all world states $\mathbf{\Omega}$. For a team of fixed size n, Proposition 2 applies across all world states. Hence, the team's size must be a multiple of a divisor (> 1) of each $|\mathbf{Good}_\omega|$. Note that the pdfs of the agents (and also r) may change according to ω. Let $\mathbf{D}$ be a set containing one divisor of each world state (if two or more world states have a common divisor x, it will be representable by only one $x \in \mathbf{D}$). Hence, $\forall \omega$, $\exists d \in \mathbf{D}$, such that $d \bigm| |\mathbf{Good}_\omega|$; and $\forall d \in \mathbf{D}$, $\exists \mathbf{Good}_\omega$, such that $d \bigm| |\mathbf{Good}_\omega|$. There are multiple possible $\mathbf{D}$ sets, from the superset of all possibilities $\mathscr{D}$.

Therefore, we can now study the minimum size necessary for an optimal team. Applying Proposition 2 at each world state ω, we have that the minimum size necessary for an optimal team is $n = \min_{\mathbf{D} \in \mathscr{D}} \prod_{d \in \mathbf{D}} d$. Hence, our worst case is when each $|\mathbf{Good}_\omega|$ is a unique prime, as the team will have to be a product of each (unique) optimal action space sizes. This means that:

Proposition 3. *In the worst case, the minimum team size is exponential in the size of the world states $|\mathbf{\Omega}|$. In the best case, the minimum necessary team size is a constant with $|\mathbf{\Omega}|$.*

Proof. In the worst case, each added world state ω has a unique prime optimal action space size. Hence, the minimum team size is at least the product of the first $|\mathbf{\Omega}|$ primes, which, by the *prime number theorem*, has growth rate $\exp((1 + o(1))|\mathbf{\Omega}| \log |\mathbf{\Omega}|)$. In the best case, each added $\mathbf{Good}_\omega$ has a common divisor with previous ones, and the minimum necessary team size does not change. ■

However, we show that the worst case happens with low probability, and the best case with high probability. Let N be the maximum possible $|\mathbf{Good}|$, and assume that each new world state ω_j will have a uniformly randomly drawn number of optimal actions, denoted as m_j, for all $j = 1, \ldots, M$.

Proposition 4. *The probability that the minimum necessary team size grows exponentially tends to 0, and the probability that it is constant tends to 1, as $M, N \to \infty$.*

Proof. We need to show that the probability that $m_1, \ldots, m_{M-1}$ are all co-prime with m_M tends to 0 as $M, N \to \infty$. Assume $N \to \infty$, then given any prime p, the probability that at least one of any independently randomly generated $M-1$ numbers $m_1, ..., m_{M-1}$ has factor p is $1 - (1 - \frac{1}{p})^{M-1}$, while the probability that one independently randomly generated number m_M has factor p is $\frac{1}{p}$. Therefore, the probability m_M shares common factor p with at least one of $m_1, \ldots, m_{M-1}$ is $\frac{1-(1-\frac{1}{p})^{M-1}}{p}$. The probability that m_M is co-prime with all $m_1, \ldots, m_{M-1}$ is $\prod_{\text{all primes } p}[1 - \frac{1-(1-\frac{1}{p})^{M-1}}{p}]$; which tends to $\prod_{\text{all primes } p}(1 - \frac{1}{p}) = \frac{1}{\zeta(1)} = 0$, where $\zeta(s)$ is the Riemann zeta function, $\zeta(1) = \prod_{\text{all primes } p} \frac{1}{1-p^{-1}} = \sum_{i=1}^{\infty} \frac{1}{i} \to \infty$ (as shown by Euler). Hence, with high probability, when adding a new world state ω_j, $|\mathbf{Good}_{\omega_\mathbf{j}}|$ will share a common factor with a world state already in $\mathbf{\Omega}$. ■

Finally, we show that a diverse team of agents is always optimal as the team grows, as long as it grows *carefully*:

Theorem 1. *Let $\mathbf{D} \in \mathscr{D}$ be a set containing one factor from each $\mathbf{Good}_\omega$. For arbitrary n, the probability that we can generate an optimal team of size n converges to 0 as $|\mathbf{\Omega}| \to \infty$. However, if $n = c\prod_{d\in\mathbf{D}} d$, then the probability that the team is optimal tends to 1 as $c \to \infty$.*

Proof. For arbitrary n, let $\mathbf{P}$ be the set of its prime factors. Given one $p \in \mathbf{P}$, the probability that p is not a factor of $|\mathbf{Good}_\omega|$ is $1 - 1/p$. The probability that all $p \in \mathbf{P}$ are not factors is: $\prod_p(1-1/p)$. As $0 < \prod_p(1-1/p) < 1$, the probability that at least one $p \in \mathbf{P}$ is a factor of $|\mathbf{Good}_\omega|$ is $1-\prod_p(1-1/p) < 1$. For $|\mathbf{\Omega}|$ tests, the probability that at least one p is a factor in all of them is: $(1-\prod_p(1-1/p))^{|\mathbf{\Omega}|}$, which $\to 0$, as $|\mathbf{\Omega}| \to \infty$. Hence, the probability that $\gcd(n, |\mathbf{Good}_\omega|) = 1$ for at least one ω tends to 1, and the probability that the team can be optimal tends to 0. However, if $n = c\prod_{d\in\mathbf{D}} d$, then $gcd(n, |\mathbf{Good}_\omega|) \neq 1\ \forall\omega \in \mathbf{\Omega}$. Let N_j be the number of agents ϕ_i that have a_j in its $\mathbf{Good}^i$. As each a_j has equal probability of being selected to be in a $\mathbf{Good}^i$, for a large number of drawings ($c \to \infty$), $P(N_i = N_j) \to 1, \forall a_i, a_j \in \mathbf{Good}_\omega, \forall\omega$ (*law of large numbers*). ■

If it is expensive to test values for n such that Theorem 1 is satisfied, we can choose $n = c\prod_\omega |\mathbf{Good}_\omega|$, as it immediately follows the conditions of the theorem. Moreover, if we know the size of all $|\mathbf{Good}_\omega|$, we can check if n and $|\mathbf{Good}_\omega|$ are co-prime in $O(h)$ time (where h is the number of digits in the smaller number), using the *Euclidean algorithm*. Hence, we can test all world states in $O(|\mathbf{\Omega}|h)$ time.

4.1 Generalizations

We first show that Theorem 1 still applies for agents ϕ_i with different probabilities over optimal actions p_{Good^i}. We consider a more general definition of *optimal* team: the difference between the probabilities of picking each optimal action must be as small as possible; i.e., let $p_j^{\mathbf{\Phi}}$ be the probability of team $\mathbf{\Phi}$ picking optimal action a_j, the optimal team is such that $\Delta := \sum_{a_k}\sum_{a_l}|p_k^{\mathbf{\Phi}} - p_l^{\mathbf{\Phi}}|, \forall a_k, a_l \in \mathbf{Good}$ is minimized (hence in the previous case $\Delta = 0$).

Proposition 5. *Theorem 1 still applies when $|p_{Good^i} - p_{Good^j}| \leq \epsilon$, $\forall\phi_i, \phi_j$, for small enough $\epsilon > 0$.*

Proof. Let $\mathbf{\Phi}$ be an optimal team, where p_{Good^i} is the same for all agents ϕ_i. Hence, the probability of all actions in $\mathbf{Good}$ being selected by the team is the same. I.e., $p_k^{\mathbf{\Phi}} = p_l^{\mathbf{\Phi}}, \forall a_k, a_l \in \mathbf{Good}$, and $\Delta = 0$.

We prove by mathematical induction. Assume we change the p_{Good^i} of x agents ϕ_i, and Δ is as small as possible. Now we will change $x+1$ agents. Let's pick one agent ϕ_i and increase its p_{Good^i} by $\delta \leq \epsilon$. It follows that $p_k^{\mathbf{\Phi}} > p_l^{\mathbf{\Phi}}, \forall a_k \in \mathbf{Good^i}, a_l \notin \mathbf{Good^i}$, and the new $\Delta' := \sum_{a_k\in\mathbf{Good}}\sum_{a_l\in\mathbf{Good}}|p_k^{\mathbf{\Phi}} - p_l^{\mathbf{\Phi}}| > \Delta$.

If we add one more agent ϕ_j, such that $\mathbf{Good^j} \cap \mathbf{Good^i} = \emptyset$, the probability of voting for actions $a_m \in \mathbf{Good^j}$ increases. For small enough ϵ, p_{Good^j} will be too large to precisely equalize the probabilities, and it follows that

$p_m^{\mathbf{\Phi}} > p_k^{\mathbf{\Phi}} > p_l^{\mathbf{\Phi}}, \forall a_m \in \mathbf{Good^j}, a_k \in \mathbf{Good^i}, a_l \notin \mathbf{Good^i} \cup \mathbf{Good^j}$, and $\Delta'' := \sum_{a_k \in \mathbf{Good}} \sum_{a_l \in \mathbf{Good}} |p_k^{\mathbf{\Phi}} - p_l^{\mathbf{\Phi}}| > \Delta'$. The same applies for each newly added agent, until we have a new team such that $n = c \prod_{d \in \mathbf{D}} d$ (again, satisfying the conditions of the theorem).

The base case follows trivially. If we did not change the probability of any agent (i.e., $x = 0$), and we now increase p_{Good^i} of a single agent ϕ_i, $p_k^{\mathbf{\Phi}} > p_l^{\mathbf{\Phi}}, \forall a_k \in \mathbf{Good^i}, a_l \notin \mathbf{Good^i}$, and $\Delta' > \Delta$. By the same argument as before, adding more agents will only increase Δ', until $n = c \prod_{d \in \mathbf{D}} d$. ■

We also generalize to the case where the number of preferred actions r changes for each agent. Let the number of actions in the $\mathbf{Good^i}$ of agent ϕ_i (r^i) be decided according to a uniform distribution in the interval $[1, r']$.

Proposition 6. *If $n = r' \times c \prod_{d \in \mathbf{D}} d$, the probability that the team is optimal $\to 1$ as $c \to \infty$.*

Proof. For large n, the number of agents with $r^i = 1, \ldots, r'$ is the same. Therefore, if for each subset $\mathbf{\Phi^i} \subset \mathbf{\Phi}$, such that $r^\phi = i, \forall \phi \in \mathbf{\Phi^i}$, we have that $p_k^{\mathbf{\Phi^i}} = p_l^{\mathbf{\Phi^i}}$, $\forall a_k, a_l \in \mathbf{Good}$, we will have that $p_k^{\mathbf{\Phi}} = p_l^{\mathbf{\Phi}}, \forall a_k, a_l \in \mathbf{Good}$. Given an optimal team of size n, we have r' subsets $\mathbf{\Phi^i}$ of size n/r' each. It follows by Theorem 1 that $n/r' = c \prod_{d \in \mathbf{D}} d$, and $n = r' \times n/r' = r' \times c \prod_{d \in \mathbf{D}} d$. ■

In the next section we perform experiments with agents whose pdfs differ, and diverse teams still significantly outperform uniform teams.

5 Experiments

We run experiments with diverse and uniform teams (henceforth *diverse* and *uniform*). First, we run synthetic experiments, where we randomly create pdfs for the agents, and simulate voting iterations across a series of world states. We repeat all our experiments 100 times, and in the graphs we plot the average and the confidence interval of our results (with $p = 0.01$). We run 1000 voting iterations (z), and measure how many optimal solutions the team is able to find. We study a scenario where the number of actions ($|\mathbf{A}|$) = 100, and the number of optimal actions per world state ($|\mathbf{Good}_\omega|$) is, respectively: $< 2, 3, 5, 5, 5 >$, in a total of 750 optimal solutions.

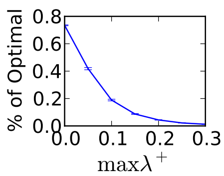

Fig. 1: Percentage as max $\boldsymbol{\lambda}^+$ grows.

At each repetition of our experiment, we randomly create a pdf for the agents. We start by studying the impact of max $\boldsymbol{\lambda}^+$ in *uniform*. When creating the *uniform* team, the total probability of playing any of the optimal actions (i.e., p_{Good}) is randomly assigned (uniform distribution) between 0.6 and 0.8. We fix the size of the team (25) and evaluate different max $\boldsymbol{\lambda}^+$ (Figure 1). As expected from Proposition 1, for max $\boldsymbol{\lambda}^+ = 0$ the system finds the highest number of optimal solutions; and as max $\boldsymbol{\lambda}^+$ increases, it quickly drops.

We then study the impact of increasing the number of agents, for *uniform* and *diverse*. To generate a *diverse* team, we draw randomly a r_ω in an interval U for each world state, that will be the size of $|\mathbf{Good}^i|$. We study three variants: *diverse**, where $U = (0, |\mathbf{Good}_\omega|]$; *diverse*, where $U = (0, |\mathbf{Good}_\omega|)$, and *diverseΔ*, where we allow agents to have different r^i_ω, also drawn from $(0, |\mathbf{Good}_\omega|)$. We independently create pdfs randomly for each agent ϕ_i. For each agent we draw a number between 0.6 and 0.8 to distribute over the set of optimal actions, and randomly decide r_ω actions to compose its $\mathbf{Good}^i$ set. We distribute equally 80% of the probability of voting over optimal actions on the actions of that set.

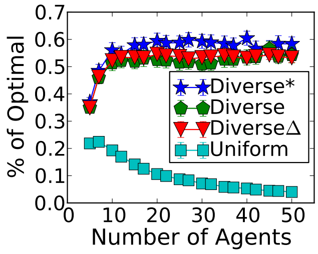

Fig. 2: Percentage of optimal solutions as # agents grows.

As we can see (Figure 2), the number of solutions *decreases* for *uniform* as the number of agents grows. Normally, in social choice, we expect the performance to improve, so this is a novel result. It is, however, expected from our Proposition 1. *Diverse*, on the other hand, improves in performance for all 3 versions, as predicted by our theory. However, the system seems to converge for a fixed z, as the performance does not increase much after around 20 agents. Hence, in Figure 3 we study larger *diverse* (continuous line) and *diverseΔ* teams (dashed line), going all the way up to 1800 agents. We also study four different number of voting iterations (z, shown in the figure by different lines): 1000, 2000, 3000, 4000. As we can see, although adding more agents was not really improving the performance in the experimental scenario under study, there is clearly a statistically significant improvement by increasing the number of voting iterations, with the system improving from around 53% of the optimal solutions, all the way up to finding more than 80% of them. However, there is a *diminishing returns* effect, as the impact of adding more iterations decreases as the actual number of iterations grows larger. We also note that *diverseΔ* is better than *diverse*, and the difference increases as z grows.

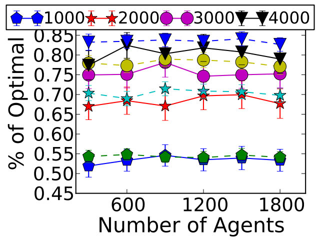

Fig. 3: Percentage for larger teams.

5.1 Experiments in Architecture Design

We study a real system for architectural building design. This is a fundamental domain, since the design of a building impacts its energy usage during its whole life-span [2, 13]. We use Beagle [8], a multi-objective design optimization software that assists users in the early stage design of buildings. Hence, the experiments presented here were run in an actual system, that performs expensive energy evaluations over complex architectural designs, and represent months of experimental work.

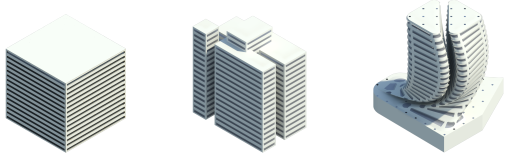

(a) Base (b) Office Park (c) Contemporary

Fig. 4: Parametric designs with increasing complexity used in our experiments.

First, the designer creates a parametric design, containing (as discussed in Section 3) a set of parameters that can be modified within a specified range, allowing the creation of many variations. We use designs from Gerber and Lin (2013) [8]: *base*, a simple building type with uniform program (i.e., tenant type); *office park*, a multi-tenant grouping of towers; and *contemporary*, a double "twisted" tower that includes multiple occupancy types, relevant to contemporary architectural practices. We show the designs in Figure 4.

Beagle uses a Genetic Algorithm (GA) to optimize the building design based on three objectives: energy efficiency, financial performance and area requirements. In detail, the objective functions are: S_{obj} : $\max SPCS$; E_{obj} : $\min EUI$; F_{obj} : $\max NPV$. SPCS is the Spatial Programming Compliance Score, EUI is the Energy Use Intensity and NPV is the Net Present Value, defined as follows.

SPCS defines how well a building conforms to the project requirements (by measuring how close the area dedicated to different activities is to a given specification). Let **L** be a list of activities (in our designs, **L**=<Office, Hotel, Retail, Parking>), $area(l)$ be the total area in a building dedicated to activity l and $requirement(l)$ be the area for activity l given in a project specification. SPCS is defined as: $SPCS = 100 * \left(1 - \frac{\sum_{l \in \mathbf{L}} |area(l) - requirement(l)|}{|\mathbf{L}|}\right)$

EUI regulates the overall energy performance of the building. This is an estimated overall building energy consumption in relation to the overall building floor area. The process to obtain the energy analysis result is automated in Beagle through Autodesk Green Building Studio (GBS) web service.

Finally, **NPV** is a commonly used financial evaluation. It measures the financial performance for the whole building life cycle, given by: $NPV = \left(\sum_{t=1}^{T} \frac{c_t}{(1+r)^t}\right) - c_0$, where T is the Cash Flow Time Span, r is the Annual Rate of Return, c_0 is the construction cost, and c_t = Revenue − Operation Cost.

Many options affect the execution of the GA, including: initial population size, size of the population, selection size, crossover ratio, mutation ratio, maximum iteration. Further details about Beagle are at Gerber and Lin (2013) [8].

In the end of the optimization process, the GA outputs a set of solutions. These are considered "optimal", according to the internal evaluation of the GA, but are not necessarily so. As in our theory, for each parameter the assigned value is going to be one of the optimal ones with a certain probability. In fact, most of the solutions outputted by the GAs are later identified as sub-optimal and eliminated in comparison with better ones found by the teams.

We model each run of the GA as an agent ϕ. Each parameter of the parametric design is a world state ω, where the agents decide among different actions $\mathbf{A}$ (i.e., possible values for the current parameter). Our model assumes independent multiple voting iterations across all world states. However, as in general it could be expensive to pool agents for votes in a large number of iterations, we test a more realistic scenario by pooling only 3 solutions per agent, but running multiple voting iterations by aggregating over all possible combinations of them, in a total of 81 voting iterations.

Agent	PZ	SZ	CR	MR
Agent 1	12	10	0.8	0.1
Agent 2	18	8	0.6	0.2
Agent 3	24	16	0.55	0.15
Agent 4	30	20	0.4	0.25

Table 1: GA parameters for the diverse team. Initial Population and Maximum Iteration were kept as constants: 10 and 5, respectively. PZ = Population Size, SZ = Selection Size, CR = Crossover Ratio, MR = Mutation Ratio.

We create 4 different agents, using different options for the GA (as shown in Table 1). Contrary to the previous synthetic experiments, we are dealing here with real (and consequently complex) design problems. Hence, the true set of optimal solutions is unknown. We approach the problem in a comparative fashion: when evaluating different systems, we consider the union of the set of solutions of all of them. That is, let $\mathbf{H}_x$ be the set of solutions of system x; we consider the set $\mathcal{H} = \bigcup_x \mathbf{H}_x$. We compare all solutions in $\mathcal{H}$, and consider as optimal the best solutions in $\mathcal{H}$, forming the set of optimal solutions $\mathcal{O}$. We use the concept of pareto dominance: the best solutions in $\mathcal{H}$ are the ones that *dominate* all other solutions (i.e., they are better in all 3 factors). As we know which system generated each solution $o \in \mathcal{O}$, we estimate the set of optimal solutions $\mathbf{S_x}$ of each system.

Although our theory focuses on *plurality voting* as the aggregation methodology, we also present results using the *mean* and the *median* of the opinions of the agents. That is, given one combination (a set of one solution from each agent), we also generate a new solution by calculating the *mean/median* across all parameters.

Concerning *uniform*, we evaluate a team composed of copies of the "best" agent. By "best", we mean the agent that finds the highest number of optimal solutions. According to Proposition 1, such an agent should be the one with the lowest $\max \boldsymbol{\lambda}^{+}$, and we can predict that voting among copies of that agent generates a large number of optimal actions. Hence, for each design, we first compare all solutions of all agents (i.e., construct $\mathcal{H}$ as the union of the solutions of all agents), to estimate which one has the largest set of optimal solutions $\mathbf{S}$. We, then, run that agent multiple times, creating *uniform*. For *diverse*, we consider one copy of each agent.

We aggregate the solutions of *diverse* and *uniform*. We run 81 aggregation iterations (across all parameters/world states), by selecting 3 solutions from each agent ϕ_i, in its set of solutions $\mathbf{H}_i$, and aggregating all possible combinations of these solutions. We evaluate together the solutions of all agents and all teams (i.e., we construct $\mathcal{H}$ with the solutions of all systems), in order to estimate the size of $\mathbf{S_x}$ of each system.

In Figure 5 (a), we show the percentage of optimal solutions for all systems, in relation to $|\mathcal{O}|$. For clarity, we represent the result of the individual agents by the one that had the highest percentage. As we can see, in all parametric designs the teams find a significantly larger percentage of optimal solutions than the individual agents. The agents find less than 1% of the solutions, while the teams are in general always close to or above 15%. In total (considering all aggregation methods and all agents), for all three parametric designs the agents find only about 1% of the optimal solutions, while *uniform* finds around 51% and *diverse* 47%. Looking at *vote*, in *base diverse* finds a larger percentage of optimal solutions than *uniform* (around 9.4% for *uniform*, while 11.6% for *diverse*). In *office park* and *contemporary*, however, *uniform* finds more solutions than *diverse*. Based on Proposition 1, we expect that this is caused by the best agent having a lower $\max \boldsymbol{\lambda}^{+}$ in *office park* and *contemporary* than in *base*.

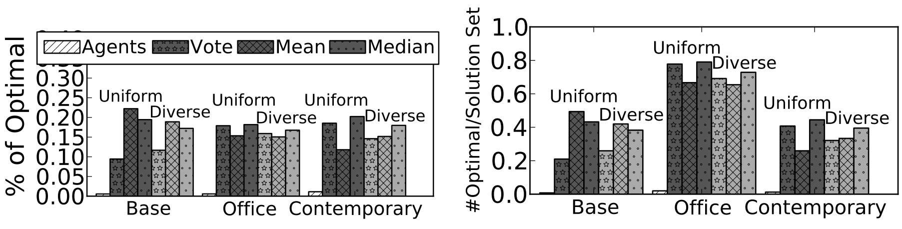

(a) Percentage in relation to all solutions found by all systems.

(b) Percentage in relation to the # of solutions of each system.

Fig. 5: Percentage of optimal solutions of each system.

Figure 5 (b) shows the percentage of optimal solutions found, in relation to the size of the set of evaluated solutions of each system. That is, let $\mathbf{O}_x$ be the set of optimal solutions of system x, in $\mathcal{O}$. We show $\frac{|\mathbf{O}_x|}{|\mathbf{H}_x|}$. Concerning *vote*, the teams are able to find a new optimal solution around 20% of the time for *base*, around

73% of the time for *office park* and around 36% of the time for *contemporary*. Meanwhile, for the individual agents it is close to 0%. We can see that teams have a great potential in generating new optimal solutions, as expected from our theory. However, as studied in our synthetic experiments, we can expect some *diminishing returns* when increasing the number of voting iterations. We show examples of solutions created by the teams in Figure 6.

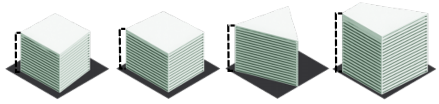

(a) Base. Shaded area shows variance of building's footprint in relation to site. Dashed line indicates height variance.

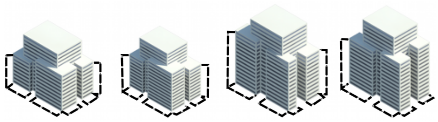

(b) Office Park. Dashed line shows variance in volume.

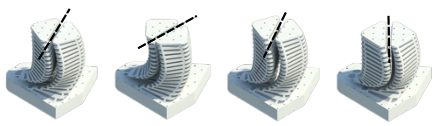

(c) Contemporary. Line shows variance in orientation.

Fig. 6: Some building designs generated by the teams.

We also plot in Figure 7 (a) the percentage of solutions that were reported to be optimal by each agent, but were later discovered to be suboptimal by evaluating $\mathcal{H}$. A large amount of solutions are eliminated (close to 100%), helping the designer to avoid making a poor decision, and increasing her confidence that the set of optimal solutions found represent well the "true" pareto frontier. Moreover, we test for duplicated solutions across different aggregation methods, different teams and different agents. The number is small: only 4 in *contemporary*, and none in *base* and *office park*. Hence, we are providing a high coverage of the pareto frontier for the designer. We show the total number of optimal solutions in Figure 7 (b). Finally, to better study the solutions proposed by the agents and teams, we plot all the optimal solutions in the factors space in Figure 8, where we show that the solutions give a good coverage of the pareto frontier.

6 Conclusion

Design imposes a novel problem to social choice: maximize the number of optimal solutions. We present a novel model for agent teams, that shows the potential of a system of voting agents to be *creative*, by generating a large number of optimal solutions to the designer. Our analysis, which builds a new connection with *number theory*, presents several novel results: (i) uniform teams are in general

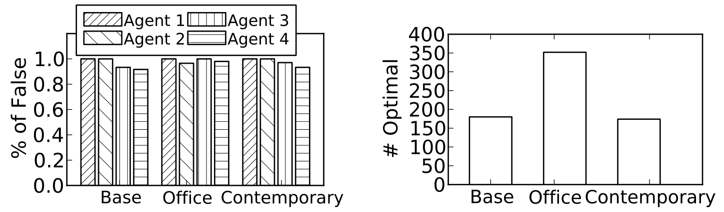

(a) False optimal solutions that are eliminated.

(b) Number of unique optimal solutions.

Fig. 7: Additional analysis.

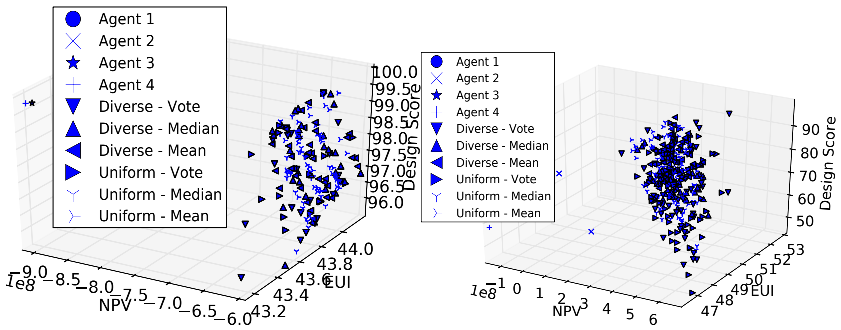

(a) Base

(b) Office Park

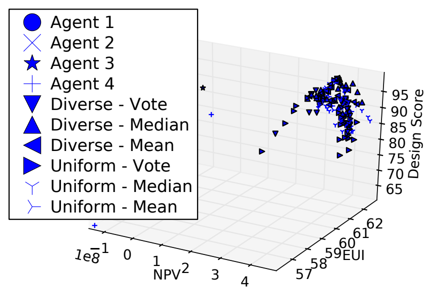

(c) Contemporary

Fig. 8: All the optimal solutions in the factor space.

suboptimal, and converge to a unique solution; (ii) diverse teams are optimal as long as the team's size grows *carefully*; (iii) the minimum optimal team size is constant with high probability; (iv) the worst case for teams is a prime number of optimal actions. Our experiments consider bounded time and relaxed assumptions, and diverse teams still perform well. We show results in architecture, where teams find a large number of solutions for designing energy-efficient buildings.

Acknowledgments: This research is supported by MURI grant W911NF-11-1-0332, and the National Science Foundation under grant 1231001.

References

1. Bachrach, Y., Graepel, T., Kasneci, G., Kosinski, M., Van Gael, J.: Crowd IQ: aggregating opinions to boost performance. In: AAMAS. pp. 535–542 (2012)
2. Bogensttter, U.: Prediction and optimization of life-cycle costs in early design. Building, Research & Information 28, 376–386 (2000)
3. Boutilier, C.: A POMDP formulation of preference elicitation problems. In: AAAI (2002)
4. Caragiannis, I., Procaccia, A.D., Shah, N.: When do noisy votes reveal the truth? In: EC. pp. 143–160. ACM, New York, NY, USA (2013)
5. Conitzer, V., Sandholm, T.: Common voting rules as maximum likelihood estimators. In: UAI'05. pp. 145–152 (2005)
6. Elkind, E., Shah, N.: Electing the most probable without eliminating the irrational: Voting over intransitive domains. In: UAI (2014)
7. Erhan, H., Wang, I., Shireen, N.: Interacting with thousands: A parametric-space exploration method in generative design. In: Proceedings of the 2014 Conference of the Association for Computer Aided Design in Architecture. ACADIA (2014)
8. Gerber, D.J., Lin, S.H.E.: Designing in complexity: Simulation, integration, and multidisciplinary design optimization for architecture. Simulation (Apr 2013)
9. Globa, A., Donn, M., Moloney, J.: Abstraction versus cased-based: A comparative study of two approaches to support parametric design. In: ACADIA (2014)
10. Isa, I.S., Omar, S., Saad, Z., Noor, N., Osman, M.: Weather forecasting using photovoltaic system and neural network. In: Proceedings of the 2010 Second International Conference on Computational Intelligence, Communication Systems and Networks. pp. 96–100. CICSyN (July 2010)
11. Jiang, A.X., Marcolino, L.S., Procaccia, A.D., Sandholm, T., Shah, N., Tambe, M.: Diverse randomized agents vote to win. In: NIPS (2014)
12. Lin, S.h.E., Gerber, D.J.: Designing-In Performance : A Framework for Evolutionary Energy Performance Feedback in Early Stage Design. Automation and Construction 38, 59–73 (2012)
13. Lin, S.H.E., Gerber, D.J.: Evolutionary energy performance feedback for design: Multidisciplinary design optimization and performance boundaries for design decision support. Energy and Buildings 84, 426–441 (2014)
14. List, C., Goodin, R.E.: Epistemic democracy: Generalizing the Condorcet Jury Theorem. Journal of Political Philosophy 9, 277–306 (2001)
15. Mao, A., Procaccia, A.D., Chen, Y.: Better Human Computation Through Principled Voting. In: AAAI (2013)
16. Marcolino, L.S., Xu, H., Jiang, A.X., Tambe, M., Bowring, E.: Give a hard problem to a diverse team: Exploring large action spaces. In: AAAI (2014)
17. Nurmi, H.: Comparing Voting Systems. Springer (1987)
18. Polikar, R.: Ensemble learning. In: Zhang, C., Ma, Y. (eds.) Ensemble Machine Learning: Methods and Applications. Springer (2012)
19. Procaccia, A.D., Reddi, S.J., Shah, N.: A maximum likelihood approach for selecting sets of alternatives. In: UAI (2012)
20. Soufiani, H.A., Parkes, D.C., Xia, L.: Random utility theory for social choice. In: NIPS. pp. 126–134 (2012)
21. Vierlinger, R., Bollinger, K.: Accommodating change in parametric design. In: ACADIA (2014)
22. Xia, L., Conitzer, V.: A maximum likelihood approach towards aggregating partial orders. In: IJCAI (2011)