Optimization for Semantic-Aware Resource Allocation under CPT-based Utilities

Symeon Vaidanis*, Photios A. Stavrou*, and Marios Kountouris*†

*Communication Systems Department, EURECOM, Sophia-Antipolis, France [†]Department of Computer Science and Artificial Intelligence, University of Granada, Spain Emails: {symeon.vaidanis, fotios.stavrou, marios.kountouris}@eurecom.fr

Abstract—The problem of resource allocation in goal-oriented semantic communication with semantic-aware utilities and subjective risk perception is studied here. By linking information importance to risk aversion, we model agent behavior using Cumulative Prospect Theory (CPT), which incorporates risk-sensitive utility functions and nonlinear transformations of distributions, reflecting subjective perceptions of gains and losses. The objective is to maximize the aggregate utility across multiple CPT-modeled agents, which leads to a nonconvex, nonsmooth optimization problem. To efficiently solve this challenging problem, we propose a new algorithmic framework that combines successive convex approximation (SCA) with the projected subgradient method and Lagrangian relaxation, Our approach enables tractable optimization while preserving solution quality, offering both theoretical rigor and practical effectiveness in semantics-aware resource allocation.

Index Terms—Goal-oriented semantic communication, resource allocation, cumulative prospect theory, risk aversion.

I. INTRODUCTION

Goal-oriented semantic communication has emerged as a transformative paradigm for future network design, focusing the entire communication process on the relevance and importance of information content, defined by application requirements and perceived by end users [1]–[3]. This framework minimizes redundant data transmission, enhances resource and computational efficiency, and serves as a cornerstone for enabling collaborative, hyperconnected intelligence and the emerging Internet of Agents.

A key challenge in semantic communication is to establish an operational and meaningful definition of information significance. In this context, messages are often subjectively evaluated and events interpreted on the basis of conceptual relevance and the decision maker's background knowledge.

To address this, we link the semantic importance of content with user behavior under risk, irrational biases in decisionmaking, and their subjective perception of event occurrence probabilities. To this end, we leverage concepts from cumulative prospect theory (CPT) [4], highlighting its potential for application in multi-objective optimization and semanticsaware resource allocation problems.

One of the main challenges in applying CPT lies in optimizing its utility function, which is inherently nonconvex and nonsmooth, even in its baseline formulation.

Recent studies have primarily addressed this challenge in the context of portfolio optimization. In [5], the authors employ

a minorization-maximization (MM) algorithm to maximize the CPT utility, showing that the problem can be approached using a convex-concave procedure that iteratively optimizes a local approximation. However, the proposed methods are heuristic and lack formal theoretical guarantees. In [6], the authors adopt the Alternating Direction Method of Multipliers (ADMM) in two distinct approaches: one combined with dynamic programming, and the other with the Pool Adjacent Violators (PAV) algorithm. While effective in practice, these approaches, particularly the PAV-based method, fail to fully account for the inherent nonsmoothness and nonconvexity of the CPT utility function.

In this work, we propose an optimization framework for solving resource allocation problems in semantic communication involving CPT-based agents. We specifically apply our method to power allocation scenarios, as considered in [7]. The core contribution lies in explicitly addressing the non-smoothness and non-convexity of the utility function during the optimization process. To this end, we integrate the Successive Convex Approximation (SCA) [8]–[11], the Projected Subgradient Method, and the Lagrangian relaxation [12]. Numerical results demonstrate that the proposed method outperforms Sequential Quadratic Programming (as implemented in MATLAB's Optimization Toolbox) in terms of scalability with the number of agents, while incurring a slightly higher, yet comparable, execution time.

II. CPT PRELIMINARIES

In this section, we provide a brief overview of the mathematical framework of CPT [13], introducing its key features absent in Expected Utility Theory (EUT).

First, agents evaluate events relative to a *reference point* x_0 , which may represent an expected or previously achieved operating level and can vary across different application scenarios. Deviations from x_0 reflect the agent's subjective perception of value and form the basis for expressing preferences and decision biases, thus capturing the semantic relevance of the outcomes.

Second, the CPT utility function u(x) exhibits *loss aversion*, meaning that agents are more sensitive to losses than to equivalent gains. Formally, this is captured by the condition $u'(x_0^+) < u'(x_0^-)$, indicating a steeper slope on the loss side of the x_0 . Symmetric bet aversion [14] is defined by $u(x_0 + \delta) + u(x_0 - \delta) < 0, \forall \delta > 0$, with $u(x_0) = 0$. This condition implies that agents reject all fair symmetric gambles in favor of the status quo. A stronger form of this aversion requires $u(x_0 + \delta_1) + u(x_0 - \delta_1) < u(x_0 + \delta_2) + u(x_0 - \delta_2), \forall 0 < \delta_2 \leq \delta_1$, indicating a growing reluctance toward larger symmetric risks. The curvature of u(x) reflects both risk attitudes and sensitivity to change within each domain. In the gain domain, a concave function implies greater sensitivity to small gains near the reference point, indicating risk aversion, while a convex function suggests increased responsiveness to larger, distant gains, implying risk seeking. The opposite holds in the loss domain. This contrasts with EUT, which typically assumes a globally concave utility function.

Third, CPT introduces *nonlinear probability distortion*, where objective probabilities are transformed through a probability weighting function (PWF) to reflect subjective perception. The PWF captures how agents tend to overweight lowprobability events and underweight moderate to high probabilities, modeling human biases in uncertainty assessment more accurately than traditional linear probability interpretations.

In this work, we adopt a recently proposed generalized formulation of CPT utility functions [7], which allows the modeling of a wider spectrum of risk attitudes and behavioral patterns.

$$u(x) = \begin{cases} \lambda_1 \frac{\mu_1 - \exp\left(\frac{\alpha}{\gamma_1} \cdot \frac{(x - x_0)}{m}\right)}{\alpha} & x_0 \le x \\ \lambda_2 \frac{\mu_2 - \exp\left(\frac{\beta}{\gamma_2} \cdot \frac{(x - x_0)}{n}\right)}{\beta} & x < x_0 \end{cases}$$
(1)

where α , β , λ_1 , λ_2 , γ_1 , γ_2 , μ_1 , μ_2 , m and n are user specific parameters generally defined on \mathbb{R} .

	Gain	Loss			
Constant	$egin{array}{l} \gamma_1 ightarrow 0^-, 0 < lpha, \ 0 < \lambda_1 \cdot \mu_1, \ 0 < m \end{array}$	$\begin{array}{c c} \gamma_1 \to 0^-, 0 < \alpha, & \gamma_2 \to 0^+, 0 < \beta, \\ < \lambda_1 \cdot \mu_1, \ 0 < m & \lambda_2 \cdot \mu_2 < 0, 0 < n \end{array}$			
Linear	$\alpha \to 0, \frac{\lambda_1}{\gamma_1} < 0, \ 0 < m$	$\beta \rightarrow 0, \frac{\lambda_2}{\gamma_2} < 0, 0 < n$			
Convex	$\frac{\lambda_1}{\gamma_1} < 0.0 < \frac{\alpha}{\gamma_1},$ $\mu_1 \le 1, \ 0 < m$	$\begin{array}{c} \frac{\lambda_2}{\gamma_2} < 0.0 < \frac{\beta}{\gamma_2}, \\ 1 \le \mu_2, 0 < n \end{array}$			
Concave	$\frac{\frac{\lambda_1}{\gamma_1} < 0, \frac{\alpha}{\gamma_1} < 0,}{1 \le \mu_1, 0 < m}$	$\frac{\frac{\lambda_2}{\gamma_2} < 0, \frac{\beta}{\gamma_2} < 0,}{\mu_2 \le 1, 0 < n}$			

TABLE I: Summary of parameter values and possible utility function shapes across subdomains.

III. SYSTEM MODEL AND PROBLEM FORMULATION

Consider a system in which a server communicates with a set of agents $\mathcal{N} = \{1, \ldots, N\}$ over a wireless medium. Each agent exchanges information with the server, r, with varying degrees of semantic importance depending on its specific goal. We study a problem where resources have to be allocated among different agents based on the perceived importance of information, as evaluated using metrics derived from CPT.

For clarity of exposition, we consider the total transmit power P_{total} as the resource to be allocated and model the CPT-based metric as a function of the communication channel conditions. These conditions are characterized by the signalto-noise ratio (SNR), given by SNR = $\frac{P|h|^2}{\sigma^2}$, where h denotes the channel coefficient, P is the transmit power and σ^2 is the variance of the additive white Gaussian noise (AWGN).

Each agent employs its own parametrized utility function, and the evaluation of the CPT-based metric under a given resource allocation is independent across agents, rendering the optimization problem separable. The total resource budget is finite, imposing an upper bound on the sum of allocated resources. In the most general case, the utility of each agent is weighted by a subjectively perceived probability, e.g., the likelihood that the *i*-th agent is active. As a result, the optimization problem is separable, but inherently non-concave and non-smooth, and can be formulated as follows:

$$\min_{\mathbf{P}} \quad f(\mathbf{P})$$
s.t. $\mathbf{P} \in \mathcal{S}$ (2)
 $g(\mathbf{P}) \leq 0$

where $f(\mathbf{P}) = -\sum_{i=1}^{N} w(p_i) f_i(P(i))$, $f_i(P(i)) = u_i(\text{SNR}(i))$, $\text{SNR}(i) = \frac{P(i) \cdot |h(i)|^2}{\sigma^2}, w(p_i)$ is the PWF, modeling the *i*-th agent's subjective assessment of probability $p_i, \mathcal{S} = \mathbb{R}^N_+$ and $g(\mathbf{P}) = \sum_{i=1}^{N} P(i) - P_{total}$.

IV. OPTIMIZATION PROBLEM

In this section, we present a framework for solving general resource allocation problems that involve CPT-based agents. Our approach combines SCA with the PSM. We first present a general solution to the CPT-based resource allocation problem, which can then be particularized to specific scenarios, such as the power allocation problem. In that case, the parameters of the utility functions can take the following values: $\alpha(i)$, $\beta(i) \in \mathbb{R}$ and $\lambda_1(i)$, $\lambda_2(i) \in \mathbb{R}^*$ and $\gamma_1(i)$, $\gamma_2(i) \in \mathbb{R}^*$ and $\mu_1(i) = \mu_2(i) = 1$ for $i \in \{1, \ldots, N\}$. The reference point of the *i*-th agent is denoted as SNR₀(*i*).

A. Successive Convex Approximation

The optimization problem in 2 can contain convex (from a maximization perspective) or linear components, depending on the values of the CPT parameters. To handle this, we employ the SCA algorithm, which iteratively solves the original problem using surrogate functions. The SCA algorithm, adapted to the notation of the optimization problem 2, is detailed in Algorithm 1. In essence, the SCA method approximates the concave (or convex) and linear parts of the objective function, with respect to whether the goal is maximization or minimization, by surrogate functions. The optimization problem is then solved iteratively, leveraging the results of the previous iteration at each step, until convergence is achieved.

The main construction rules for the surrogate function in the context of a minimization problem are as follows:

- 1) $\tilde{f}(\mathbf{x}|\mathbf{x}^{(k)})$ should be strongly convex over the domain \mathcal{X} .
- 2) $\nabla \tilde{f}(\mathbf{x}^{(k)}|\mathbf{x}^{(k)}) = \nabla f(\mathbf{x}^{(k)}), \forall \mathbf{x}^{(k)} \in \mathcal{X}$
- 3) $\nabla \tilde{f}(\mathbf{x}|\mathbf{x}^{(k)})$ must be continuous on \mathcal{X}
- 4) {Optional for stronger convergence guarantees} $\nabla \tilde{f}(\mathbf{x}|\mathbf{x}^{(k)})$ is uniformly Lipschitz continuous over \mathcal{X}

where \mathcal{X} denotes the domain of \mathbf{x} and $\tilde{f}(\mathbf{x}|\mathbf{x}^{(k)})$ is the surrogate function that approximates the original function $f(\mathbf{x})$ in the (k + 1)-th iteration. It should be emphasized that these construction rules were originally proposed in [10] for smooth but nonconvex objective functions. However, the authors argue that these results can be extended to handle nonsmooth objectives as well. Specifically, if the objective function can be decomposed as the sum of a smooth but nonconvex function $U(\mathbf{x})$ (or concave, in the case of maximization) and a nonsmooth but convex function $R(\mathbf{x})$ (or concave for maximization), then the convergence guarantees applicable to smooth functions can still hold. In our case the function U and R for each agent can be the following:

$$U_{i}(x) = u_{i}(x) + \mathbb{1}\{x \ge x_{0}\} \cdot \left(\frac{\partial u_{i}}{\partial x(i)} \Big|_{x_{0}^{-}(i)} - \frac{\partial u_{i}}{\partial x(i)} \Big|_{x_{0}^{+}(i)} \right) \cdot (x - x_{0})$$

$$B_{i}(x) = -\mathbb{1}\{x \ge x_{0}\} \cdot$$
(3a)

$$\cdot \left(\frac{\partial u_i}{\partial x(i)} \Big|_{x_0^-(i)} - \frac{\partial u_i}{\partial x(i)} \Big|_{x_0^+(i)} \right) \cdot (x - x_0).$$
 (3b)

In addition, the authors of [10] emphasize that the first construction rule, i.e., the strong convexity of the surrogate function, is the most critical. This condition significantly influences the choice of the optimization method used to solve the inner surrogate problem in 5. The remaining construction rules serve as technical requirements to ensure that the surrogate function maintains the same local first-order behavior as the original function. However, due to the technical challenge introduced by the non-smoothness at the reference point in our problem, we introduce two additional construction conditions into our methodology to better handle this issue.

5)
$$\hat{f}(\mathbf{x}|\mathbf{x}^{(k)}) \ge f(\mathbf{x}), \ \forall \mathbf{x} \in \mathcal{X}$$

6) If $x^{(k)}(i) \to x_0^{(i)}$, then $\frac{\partial \tilde{u}_i}{\partial x(i)}\Big|_{x_0^+(i)} \to \frac{\partial u_i}{\partial x(i)}\Big|_{x_0^+(i)}$ and
 $\frac{\partial \tilde{u}_i}{\partial x(i)}\Big|_{x_0^-(i)} \to \frac{\partial u_i}{\partial x(i)}\Big|_{x_0^-(i)}$

Regarding the convergence of the SCA algorithm, the authors of [10] prove that if all four surrogate function conditions are satisfied, and the step size is chosen to meet the following criteria:

$$\theta^{(l)} \in (0,1], \theta^{(l)} \xrightarrow{l \to \infty} 0, \sum_{l} \theta^{(l)} = +\infty$$
(4)

then every limit point of the sequence $\mathbf{x}^{(k)}$ is a stationary point of the original optimization problem. In other cases, if the three basic conditions are satisfied, at least one feasible point of the sequence $\mathbf{x}^{(k)}$ is a stationary point of the original optimization problem. In our case, the fourth condition does not hold due to the fact that the objective function is not smooth at the reference point.

Algorithm 1 SCA algorithm for optimization problem

Set l = 0, initialize feasible point $\mathbf{P}^{(0)} \in \mathcal{S}$, step size $\{\theta^{(l)}\} \in (0, 1]$ **repeat** $\hat{\mathbf{P}}(\mathbf{P}^{(l)}) = \arg\min \tilde{f}(\mathbf{P}|\mathbf{P}^{(l)})$ s.t. $\mathbf{P} \in \mathcal{S}$ $g(\mathbf{P}) \leq 0$ $\mathbf{P}^{(l+1)} = \mathbf{P}^{(l)} + \theta^{(l)} \cdot (\hat{\mathbf{P}}(\mathbf{P}^{(l)}) - \mathbf{P}^{(l)})$ $l \leftarrow l + 1$ **until** convergence **return** $\mathbf{P}^{(l)}$

We now proceed to apply the surrogate function construction rules within our proposed framework. The general form of the surrogate functions, tailored to the structure of our problem, is defined separately for the gain and loss subdomains. For the gain subdomain, the surrogate utility function is given by:

$$\tilde{u}_{\text{gain}}(x(i)|x^{(l)}(i)) = \lambda_{\text{gain}} \cdot \frac{1 - \exp\left(\alpha_{\text{gain}} \cdot \frac{x(i) - \mu_{\text{gain}}}{m}\right)}{\alpha_{\text{gain}}} \quad (6)$$

and for the loss subdomain, the corresponding surrogate function is:

$$\tilde{u}_{\text{loss}}(x(i)|x^{(l)}(i)) = \lambda_{\text{loss}} \cdot \frac{1 - \exp\left(\beta_{\text{loss}} \cdot \frac{x(i) - \mu_{\text{loss}}}{n}\right)}{\beta_{\text{loss}}} \quad (7)$$

More precisely, we categorize all possible cases into six distinct scenarios, based on the values of the utility function parameters. **Case 1**: if the parameters satisfy $\frac{\alpha(i)}{\gamma_1(i)} < 0$, $\frac{\beta(i)}{\gamma_2(i)} < 0$, then the utility function is strictly concave over the entire domain. In this case, the original function itself serves as a valid surrogate, as it inherently satisfies all the construction rules. Case 2: if $\text{SNR}_0(i) \leq \text{SNR}^{(l)}(i)$ and $\frac{\alpha(i)}{\gamma_1(i)} < 0, 0 \leq$ $\frac{\beta(i)}{\gamma_2(i)}$, the loss subdomain must be approximated using a concave surrogate function $\tilde{u}_{loss}(x(i)|x^{(l)}(i))$ with $\mu_{loss} = x_0(i)$, $\beta_{loss} \leq 0$, and $\lambda_{loss} = -n \cdot \frac{\partial u_i}{\partial x(i)}\Big|_{x_0^-(i)}$. The gain subdomain remains unchanged, since it is already concave. An interesting observation is that the 'best' concave approximation of a convex function, given a specific intersection point and subject to the surrogate construction rules, is the tangent line at that point. This choice minimizes the Euclidean distance between the original and surrogate functions locally. However, it is important to emphasize that if a linear surrogate function (such as the tangent) is used, the first construction rule (strong convexity) is not satisfied. As a result, the PSM cannot be applied in the inner minimization step of SCA under this choice.

Case 3: if the parameters satisfy $\text{SNR}_0(i) < \text{SNR}^{(l)}(i)$ and $0 \leq \frac{\alpha(i)}{\gamma_1(i)}, \frac{\beta(i)}{\gamma_2(i)} \in \mathbb{R}$, both the gain and loss subdomains require surrogate approximations. For the gain subdomain, we define the surrogate function $\tilde{u}_{\text{gain}}(x(i)|x^{(l)}(i))$ with $\mu_{\text{gain}} = x^{(l)}(i), \alpha_{\text{gain}} \leq 0, l_{\text{gain}} = u_i(x^{(l)}(i)), \lambda_{\text{gain}} =$
$$\begin{split} &-m \cdot \frac{\partial u_i}{\partial x(i)} \bigg|_{x^{(l)}(i)} \text{ and the loss subdomain with the follow-}\\ &\text{ing surrogate function } \tilde{u}_{\text{loss}}(x(i)|x^{(l)}(i)) \text{ with } \mu_{\text{loss}} = x_0(i)\\ &, \ \beta_{\text{loss},1} \le \min\left\{0, \frac{\beta(i)}{\gamma_2(i)}\right\}, \ |\beta_{\text{loss},2}| \le |\beta(i)| \text{ and } 0 < \\ &\beta_{\text{loss},2} \cdot \beta(i), \ l_{\text{loss}} = \tilde{u}_{\text{gain}}(x_0(i)|x^{(l)}(i)), \ \lambda_{\text{loss}} = -n \cdot \frac{\beta_{\text{loss},2}}{\beta_{\text{loss},1}} \cdot \\ &\max\left\{\frac{\partial u_i}{\partial x(i)}\bigg|_{x_0^-(i)}, \frac{\partial \tilde{u}_{\text{gain}}}{\partial x(i)}\bigg|_{x_0^+(i)}\right\}. \end{split}$$

Case 4: if $\operatorname{SNR}^{(l)}(i) \leq \operatorname{SNR}_0(i)$ and $0 \leq \frac{\alpha(i)}{\gamma_1(i)}, \frac{\beta(i)}{\gamma_2(i)} < 0$, we approximate the gain subdomain using a concave surrogate function $\tilde{u}_{\text{gain}}(x(i)|x^{(l)}(i))$ with parameters $\mu_{\text{gain}} = x_0(i)$, $\alpha_{\text{gain}} \leq 0$ and $\lambda_{\text{gain}} = -m \cdot \frac{\partial u_i}{\partial x(i)} \Big|_{x_0^+(i)}$. The loss subdomain is concave and therefore kept unchanged. **Case 5**: if $\operatorname{SNR}^{(l)}(i) < \operatorname{SNR}_0(i)$ and $\frac{\alpha(i)}{\gamma_1(i)} \in \mathbb{R}, 0 \leq \frac{\beta(i)}{\gamma_2(i)}$, both subdomains require surrogate approximations. The loss subdomain is approximated using the surrogate function $\tilde{u}_{\text{loss}}(x(i)|x^{(l)}(i))$ with $\mu_{\text{loss}} = x^{(l)}(i)$, $\beta_{\text{loss}} \leq 0$, $l_{\text{loss}} = u_i(x^{(l)}(i))$, $\lambda_{\text{gain}} = -n \cdot \frac{\partial u_i}{\partial x(i)} \Big|_{x^{(l)}(i)}$ and the gain subdomain with the following surro-

gate function $\tilde{u}_{gain}(x(i)|x^{(l)}(i))$ with $\mu_{gain} = x_0(i)$, $\alpha_{gain,1} \le 0$ if $0 \le \frac{\alpha(i)}{\gamma_1}$ or $\frac{\alpha(i)}{\gamma_1} \le \alpha_{gain,1} \le 0$ if $\frac{\alpha(i)}{\gamma_1} < 0$, $|\alpha(i)| \le |\alpha_{gain,2}|$ and $0 < \alpha_{gain,2} \cdot \alpha(i)$, $l_{gain} = \tilde{u}_{loss}(x_0(i)|x^{(l)}(i))$, $\lambda_{gain} = -m \cdot \frac{\alpha_{gain,2}}{\alpha_{gain,1}} \cdot \min\left\{\frac{\partial u_i}{\partial x(i)}\Big|_{x_0^+(i)}, \frac{\partial \tilde{u}_{loss}}{\partial x(i)}\Big|_{x_0^-(i)}\right\}$.

Case 6: if $\text{SNR}^{(l)}(i) = \text{SNR}_0(i)$ and $0 \le \frac{\alpha(i)}{\gamma_1(i)}, 0 \le \frac{\beta(i)}{\gamma_2(i)}$, both subdomains must be approximated using concave surrogate functions centered at the reference point. For the loss subdomain, we use the surrogate function $\tilde{u}_{\text{loss}}(x(i)|x^{(l)}(i))$ with $\mu_{\text{loss}} = x_0(i)$, $\beta_{\text{loss}} \le 0$ and $\lambda_{\text{loss}} = -n \cdot \frac{\partial u_i}{\partial x(i)} \Big|_{x_0^-(i)}$. For the gain subdomain, we define the surrogate function $\tilde{u}_{\text{gain}}(x(i)|x^{(l)}(i))$ with $\mu_{\text{gain}} = x_0(i)$, $\alpha_{\text{gain}} \le 0$ and $\lambda_{\text{gain}} = -m \cdot \frac{\partial u_i}{\partial x(i)} \Big|_{x_0^-(i)}$. This configuration ensures that the surrogate utility function is piecewise concave, follows the same first-order behavior around the reference point as the original function, and satisfies the SCA construction rules required for convergence. As both subdomains start at the same point, this case also provides a numerically stable approximation around the reference threshold.

B. Projected Subgradient Method

Regarding the optimization problem 5, we use the Lagrangian relaxation approach, following the methodology outlined in [12]. To start with, we define the Lagrangian of the problem

$$\mathcal{L}(\mathbf{P},k) = \tilde{f}\left(\mathbf{P}|\mathbf{P}^{(l)}\right) + k \cdot g\left(\mathbf{P}\right) , \ 0 \le k$$
(8)

Additionally, the Lagrangian function admits a unique minimizer over \mathbf{P} , which is denoted $\mathbf{P}^*(k)$. The dual problem then consists in maximizing this dual function with respect to the dual variables:

$$h(k) = \inf_{\mathbf{P} \in \mathcal{S}} \mathcal{L}(\mathbf{P}, k) = \tilde{f}\left(\mathbf{P}^*(k) | \mathbf{P}^{(l)}\right) + k \cdot g\left(\mathbf{P}^*(k)\right) , \ 0 \le k$$
(9)

The dual problem then consists in maximizing this dual function with respect to the dual variables:

$$\max_{k} \quad h(k)$$
s.t. $0 \le k$
(10)

Since Slater's condition holds, strong duality applies, and the primal problem can be solved by first finding the optimal dual point k^* , then recovering the primal solution as $\mathbf{P}^* = \mathbf{P}^*(k^*)$.

We solve the dual problem using the projected subgradient method [12],

$$k^{(i+1)} = \left(k^{(i)} - \zeta_i \cdot q^{(i)}\right)_+, \ q^{(i)} \in \partial(-h)(k^{(i)})$$
(11)

where $q^{(i)}$ is a subgradient with respect to k at the *i*-th iteration of the outer optimization problem with objective function h(k), and ζ_i is the step size. The step sizes are chosen to be square summable but not summable, i.e.,

$$0 \le \zeta_i, \quad \sum_{i=1}^{+\infty} \zeta_i^2 < +\infty, \quad \sum_{i=1}^{+\infty} \zeta_i = +\infty$$
 (12)

in order to ensure convergence of the method to an optimal point. This choice of step size is motivated by the fact that our objective function, which is the surrogate used in the SCA framework, is Lipschitz continuous, as it is concave on the right of x(l)(i) and concave but bounded on the left. Moreover, due to the specific form of h(k), the subdifferential of -h at any point contains only a single element, which implies that h is differentiable, and its gradient is given by:

$$\partial(-h)(k) = -g\left(\mathbf{P}^*(k)\right),\tag{13}$$

and the projected subgradient method for the dual problem has the form

$$\mathbf{P}^{(i)} = \arg\min_{\mathbf{P}\in\mathcal{S}} \left(\tilde{f}\left(\mathbf{P}|\mathbf{P}^{(l)}\right) + k^{(i)} \cdot g(\mathbf{P}) \right)$$
(14a)

$$k^{(i+1)} = \left(k^{(i)} + \zeta_i \cdot g\left(\mathbf{P}^{(i)}\right)\right). \tag{14b}$$

Since the set S is convex and the optimization problem 14a is convex, we can once again apply the projected subgradient method to solve it as follows:

$$\mathbf{P}^{(j+1)} = \Pi_{\mathcal{S}} \left(\mathbf{P}^{(j)} - \eta_j \cdot w^{(j)} \right)$$
(15)

where Π is the Euclidean projection on S

$$\Pi_{\mathcal{S}}(x_0) = \min_{x} \|x - x_0\|_2 : x \in \mathcal{S}$$
(16)

and as $S = \mathbb{R}^N_+$, the projection is equivalent to

$$\Pi_{\mathcal{S}}(x_0) = (x_0)_+. \tag{17}$$

The parameter η_j denotes the step of the method and is again chosen as *square summable but not summable*. Furthermore, $w^{(j)}$ represents the subgradient with respect to **P** at the *j*-th iteration of the inner optimization problem, whose objective function is $f(\mathbf{P}) + k^{(i)} \cdot g(\mathbf{P})$.



Fig. 1: 3D contour plot illustrating the convergence process for a scenario with N = 3 agents.

N	Mean % better than	Mean % better than	
	MATLAB (0% tol.)	MATLAB (2% tol.)	
10	38.9	49.8	
30	72.8	73.6	
50	96.2	96.4	

TABLE II: Comparison of objective function values: proposed method vs. SQP (MATLAB).

V. SIMULATION RESULTS

We begin by solving the power allocation problem in a system with three agents experiencing Rayleigh fading, where risk behavior is modeled using an S-shaped CPT utility function. This simulation illustrates the convergence behavior of our proposed optimization framework through a simple yet representative scenario. As shown in the contour plot in Figure 1, despite initializing near a local minimum, the algorithm quickly converges to a local maximum of the objective function. Next, we simulate a system under three different scenarios of the agent population: N = 10, 30, 50. Each system is simulated over 500 independent iterations. In each iteration, agents are assigned distinct CPT-based utility functions, and their channel gains are independently drawn from an exponential distribution with unit mean. The average SNR under equal power allocation is set to 7 dB.

To evaluate the performance of our proposed optimization method, we use as a metric the percentage of cases where our solution achieves an equal or better outcome, within a tolerance of k%, compared to the solution obtained by MATLAB Sequential Quadratic Programming (SQP) Global Search toolbox. Specifically, a tolerance of $\pm k\%$ implies that if the relative difference between our solution and MATLAB's falls within this range, the two are considered equivalent. The results are summarized in Table II. We also report results based on the trimmed mean in Table III.

N	Mean	1% Trimmed	2% Trimmed	5% Trimmed
		mean	mean	mean
10	-10.3613	-10.4444	-10.5075	-10.5354
30	40.6332	40.8660	41.2340	41.7988
50	81.3611	81.5843	81.6584	81.4217

TABLE III: Trimmed mean of the percentage of quantitative better solution than SQP (MATLAB).

As shown in the results tables, our proposed method outperforms MATLAB's toolbox as the number of agents increases. However, this performance gain comes with a longer execution time. Specifically, our method is significantly slower for N =10 and N = 30. Nevertheless, at N = 50, the gap narrows, with our approach being approximately four times slower. This trend suggests that as the agent population increases, our method not only maintains its performance advantage over MATLAB but also achieves comparable execution times.

ACKNOWLEDGMENT

This work is partially supported by the European Research Council (ERC) under the EU's Horizon 2020 research and innovation programme (Grant agreement No. 101003431), which partially supports the work of M. Kountouris. The work of S. Vaidanis and P. A. Stavrou is supported by the SNS JU project 6G-GOALS [2] under the EU's Horizon programme Grant Agreement No. 101139232. S. Vaidanis is also supported by the Onassis Foundation (Scholarship ID: F ZU 076-1/2024-2025).

REFERENCES

- M. Kountouris and N. Pappas, "Semantics-empowered communication for networked intelligent systems," *IEEE Communications Magazine*, vol. 59, no. 6, pp. 96–102, June 2021.
- [2] E. C. Strinati et al., "Goal-oriented and semantic communication in 6G AI-native networks: the 6G-GOALS approach," in *European Conference* on Networks and Communications & 6G Summit, June 2024.
- [3] P. A. Stavrou and M. Kountouris, "The role of fidelity in goal-oriented semantic communication: A rate distortion approach," *IEEE Trans. Commun.*, vol. 71, no. 7, pp. 3918–3931, 2023.
- [4] A. Tversky and D. Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," J. Risk Uncertain., vol. 5, no. 4, pp. 297– 323, 1992.
- [5] E. Luxenberg, P. Schiele, and S. Boyd, "Portfolio optimization with cumulative prospect theory utility via convex optimization," *Computational Economics*, vol. 64, p. 3027–3047, 2024.
- [6] X. Cui, R. Jiang, Y. Shi, R. Xiao, and Y. Yan, "Decision making under cumulative prospect theory: An alternating direction method of multipliers," *INFORMS Journal on Computing*, Sep. 2024.
- [7] S. Vaidanis, P. A. Stavrou, and M. Kountouris, "Goal-oriented semantic resource allocation with cumulative prospect theoretic agents," in *IEEE Intern. Conf. on Communications (ICC)*, 2025.
- [8] G. Scutari, F. Facchinei, P. Song, D. P. Palomar, and J.-S. Pang, "Decomposition by partial linearization: Parallel optimization of multiagent systems," *IEEE Trans. on Signal Processing*, vol. 62, no. 3, pp. 641–656, 2014.
- [9] G. Scutari and Y. Sun, "Parallel and distributed successive convex approximation methods for big-data optimization," 2018. [Online]. Available: https://arxiv.org/abs/1805.06963
- [10] G. Scutari, F. Facchinei, and L. Lampariello, "Parallel and distributed methods for constrained nonconvex optimization—part i: Theory," pp. 1929–1944, 2017.
- [11] G. Scutari, F. Facchinei, L. Lampariello, S. Sardellitti, and P. Song, "Parallel and distributed methods for constrained nonconvex optimizationpart ii: Applications in communications and machine learning," *IEEE Trans. on Signal Processing*, vol. 65, no. 8, pp. 1945–1960, 2017.
- [12] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [13] P. P. Wakker, Prospect theory: for risk and ambiguity. Cambridge University Press, 2010.
- [14] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econom.*, vol. 47, no. 2, pp. 263–291, 1979.