Feasible approximation of matching equilibria for large-scale matching for teams problems

Qikun Xiang

Nanyang Technological University, Singapore

July 5, 2024

Joint work with:

Ariel Neufeld,
Nanyang Technological University, Singapore

Matching for teams

- Matching for teams is a multi-agent game with N **populations of agents** introduced by Carlier and Ekeland [2010].
- The **types of agents** within the i-th population are described by a probability measure $\mu_i \in \mathcal{P}(\mathcal{X}_i)$ on the **type space** $\mathcal{X}_i$, where $(\mathcal{X}_i, d_{\mathcal{X}_i})$ is a compact metric space.
- There is a type of indivisible good that comes in different **qualities** described by a probability measure $\nu \in \mathcal{P}(\mathcal{Z})$ on the **quality space** $\mathcal{Z}$, where $(\mathcal{Z}, d_{\mathcal{Z}})$ is a compact metric space.
- One agent from each population must come together to form a team in order to trade a unit of good, subject to **matching costs** $c_i : \mathcal{X}_i \times \mathcal{Z} \to \mathbb{R}$, which is a continuous function.
- Each agent from the i-th population receives $\varphi_i(z)$ from trading a unit of good with quality z, where $\varphi_i : \mathcal{Z} \to \mathbb{R}$ is the **transfer function** that is continuous.
- The **matching** between the i-th population of agents and the good is described by a joint probability measure $\gamma_i \in \mathcal{P}(\mathcal{X}_i \times \mathcal{Z})$.

Matching for teams

Definition (Matching equilibrium [Carlier and Ekeland 2010])

A **matching equilibrium** $(\varphi_i)_{i=1:N}$, $(\gamma_i)_{i=1:N}$, ν satisfies:

- **(conservation)** for $i = 1, \ldots, N$, $\gamma_i \in \Gamma(\mu_i, \nu) := \big\{\gamma \in \mathcal{P}(\mathcal{X}_i \times \mathcal{Z})$: the marginals of γ on $\mathcal{X}_i$ and $\mathcal{Z}$ are μ_i and $\nu\big\}$;
- **(balance)** $\sum_{i=1}^{N} \varphi_i(z) = 0$ for all $z \in \mathcal{Z}$;
- **(rationality)** for $i = 1, \ldots, N$, $\varphi_i^{c_i}(x_i) + \varphi_i(z) = c_i(x_i, z)$ for γ_i-almost all $(x_i, z) \in \mathcal{X}_i \times \mathcal{Z}$, where

$$\varphi_i^{c_i}(x_i) := \inf_{z \in \mathcal{Z}} \big\{c_i(x_i, z) - \varphi_i(z)\big\} \qquad \forall x_i \in \mathcal{X}_i. \qquad (c_i\text{-transform of } \varphi_i)$$

- In particular, due to the **Kantorovich duality**:

$$\inf_{\gamma_i \in \Gamma(\mu_i, \nu)} \left\{ \int_{\mathcal{X}_i \times \mathcal{Z}} c_i \, \mathrm{d}\gamma_i \right\} = \sup_{\varphi_i \in \mathcal{C}(\mathcal{Z})} \left\{ \int_{\mathcal{X}_i} \varphi_i^{c_i} \, \mathrm{d}\mu_i + \int_{\mathcal{Z}} \varphi_i \, \mathrm{d}\nu \right\}.$$

the **rationality condition** implies that γ_i solves the **optimal transport** problem:

$$W_{c_i}(\mu_i, \nu) := \inf_{\gamma_i \in \Gamma(\mu_i, \nu)} \left\{ \int_{\mathcal{X}_i \times \mathcal{Z}} c_i \, \mathrm{d}\gamma_i \right\}.$$

Characterization of matching equilibrium

Theorem (Characterization of matching equilibrium [Carlier and Ekeland 2010])

1. *There exists a matching equilibrium.*

2. $(\tilde{\varphi}_i)_{i=1:N}$, $(\tilde{\gamma}_i)_{i=1:N}$, *and* $\tilde{\nu}$ *form a matching equilibrium if and only if:*
   - $\tilde{\nu}$ *is an* ***optimizer*** *of:*

$$\inf_{\nu\in\mathcal{P}(\mathcal{Z})}\left\{\sum_{i=1}^{N} W_{c_i}(\mu_i,\nu)\right\}; \qquad \text{(MT)}$$

   - *for* $i=1,\ldots,N$, $\tilde{\gamma}_i$ *is an optimizer of* $\inf_{\gamma_i\in\Gamma(\mu_i,\nu)}\left\{\int_{\mathcal{X}_i\times\mathcal{Z}} c_i \,\mathrm{d}\gamma_i\right\}$, *i.e.,* $\int_{\mathcal{X}_i\times\mathcal{Z}} c_i \,\mathrm{d}\tilde{\gamma}_i = W_{c_i}(\mu_i,\tilde{\nu})$.
   - $(\tilde{\varphi}_i)_{i=1:N}$ *is an* ***optimizer*** *of:*

$$\sup\left\{\sum_{i=1}^{N}\int_{\mathcal{X}_i}\varphi_i^{c_i}\,\mathrm{d}\mu_i : (\varphi_i:\mathcal{Z}\to\mathbb{R})_{i=1:N} \text{ are continuous}, \sum_{i=1}^{N}\varphi_i=0\right\}; \qquad \text{(MT}^*\text{)}$$

3. (MT) *and* (MT^*) *have identical optimal values.*

Example 1: business locations

- Consider a business in a city which hires $N-1$ **categories of employees** that is choosing the locations of business outlets.
  - $\mu_1, \ldots, \mu_{N-1}$: geographical distributions of employees' **dwellings**;
  - $c_1, \ldots, c_{N-1}$: employees' **commuting costs**;
  - $\varphi_1, \ldots, \varphi_{N-1}$: employees' **salary**;
  - $\gamma_1, \ldots, \gamma_{N-1}$: employees' **workplace choices**;
  - μ_N: geographical distribution of **suppliers**;
  - c_N: business's **restocking cost**;
  - φ_N: negative of business's **total salary payout**;
  - γ_N: business's **choices of outlet locations**;
  - ν: geographical distribution of **business outlets**.
- At **(matching) equilibrium**:
  - the total salary payout needs to be **balanced** with the total salary received by the employees, i.e., $-\varphi_N(z) = \sum_{i=1}^{N-1} \varphi_i(z)$ for all $z \in \mathcal{Z}$;
  - employees choose workplace **rationally** and business owners choose the business outlet locations **rationally**, i.e., for $i = 1, \ldots, N$, $\varphi_i^{c_i}(x_i) + \varphi_i(z) = c_i(x_i, z)$ for γ_i-almost all $(x_i, z) \in \mathcal{X}_i \times \mathcal{Z}$.

Example 2: Wasserstein barycenter

- For $p \in [1, \infty)$, the optimal transport problem with cost $d_{\mathcal{X}}(\,\cdot\,,\cdot\,)^p$ induces a metric $W_p(\cdot,\,\cdot)$ called the Wasserstein distance of order p on the space of probability measures, i.e.,

$$W_p(\mu, \nu) := \left(\inf_{\gamma \in \Gamma(\mu,\nu)} \left\{ \int_{\mathcal{X}\times\mathcal{X}} d_{\mathcal{X}}(x, x')^p \, \gamma(\mathrm{d}x, \mathrm{d}x') \right\} \right)^{\frac{1}{p}}.$$

- When $\mathcal{X}_1 = \cdots = \mathcal{X}_N = \mathcal{Z} \subset \mathbb{R}^d$, $c_i(\boldsymbol{x}, \boldsymbol{z}) := \lambda_i \|\boldsymbol{x} - \boldsymbol{z}\|_2^2$ for $i = 1, \ldots, N$ where $\lambda_1 > 0, \ldots, \lambda_N > 0$, $\sum_{i=1}^N \lambda_i = 1$, (MT) corresponds to:

$$\inf_{\nu \in \mathcal{P}(\mathcal{Z})} \left\{ \sum_{i=1}^N \lambda_i W_2(\mu_i, \nu)^2 \right\}.$$

- $\tilde{\nu}$ that minimizes (MT) is called a **Wasserstein barycenter** of $\mu_1, \ldots, \mu_N$ with weights $\lambda_1, \ldots, \lambda_N$ [Agueh and Carlier 2011].
- The input measures $\mu_1, \ldots, \mu_N$ can be:
  - posterior distributions of sub-samples in Bayesian inference (e.g., Srivastava, Li, Dunson [2018]);
  - histograms representing complex objects in clustering (e.g., Ye, Wu, Zhang, Li [2017]);
  - color palette distributions in color transfer (e.g., Fan, Taghvaei, Chen [2020]), etc.

Existing methods

- Existing numerical methods for matching for teams/Wasserstein barycenter:
  - assume **parametric measures** such as Gaussian: e.g., Álvarez-Esteban et al. [2016], Chewi et al. [2020];
  - assume **discrete measures** or **discretize** continuous measures: e.g., Carlier, Oberman, Oudet [2015], Benamou et al. [2015], and Anderes, Borgwardt, Miller [2016];
  - restrict the support of ν to a prespecified finite set (i.e., **fixed-support** methods): e.g., Staib et al. [2017], Claici, Chien, Solomon [2018], Dvurechenskii et al. [2018];
  - adopt **neural network parametrizations**: e.g., Fan, Taghvaei, Chen [2020], Li et al. [2020], Korotin et al. [2021].

- Our numerical method:
  - works for general cost functions $c_1, \ldots, c_N$ and general non-discrete, non-parametric $\mu_1, \ldots, \mu_N$;
  - works in a free-support setting, i.e., does not restrict the support of ν;
  - computes **feasible and approximately optimal solutions** of (MT) and (MT*);
  - computes a **sub-optimality bound** that is typically less conservative than theoretical bounds.

Parametric approximation

- Observe that

$$\sup\left\{\sum_{i=1}^N \int_{\mathcal{X}_i} \varphi_i^{c_i}\,\mathrm{d}\mu_i : (\varphi_i)_{i=1:N} \subset \mathcal{C}(\mathcal{Z}),\ \sum_{i=1}^N \varphi_i = 0\right\}$$
$$= \sup\left\{\sum_{i=1}^N \int_{\mathcal{X}_i} \psi_i\,\mathrm{d}\mu_i : \begin{array}{l} (\varphi_i)_{i=1:N} \subset \mathcal{C}(\mathcal{Z}),\ \psi_i \in \mathcal{C}(\mathcal{X}_i)\ \forall 1 \le i \le N,\ \sum_{i=1}^N \varphi_i = 0 \\ \psi_i(x) + \varphi_i(z) \le c_i(x,z)\ \forall (x,z) \in \mathcal{X}_i \times \mathcal{Z},\ \forall 1 \le i \le N \end{array}\right\}.$$

- We obtain a **parametric approximation** of (MT*) by:
  - parametrizing φ_i with basis functions $\mathcal{H} = \{h_1, \ldots, h_k\} \subset \mathcal{C}(\mathcal{Z})$: $\varphi_i = \sum_{l=1}^k w_{i,l} h_l$,
  - parametrizing ψ_i with basis functions $\mathcal{G}_i = \{g_{i,1}, \ldots, g_{i,m_i}\} \subset \mathcal{C}(\mathcal{X}_i)$: $\psi_i = y_{i,0} + \sum_{j=1}^{m_i} y_{i,j} g_{i,j}$.

$$\begin{aligned} \underset{(y_{i,0},\boldsymbol{y}_i,\boldsymbol{w}_i)}{\text{maximize}} \quad & \sum_{i=1}^N y_{i,0} + \langle \bar{\boldsymbol{g}}_i, \boldsymbol{y}_i \rangle \\ \text{subject to} \quad & y_{i,0} + \langle \boldsymbol{g}_i(x), \boldsymbol{y}_i \rangle + \langle \boldsymbol{h}(z), \boldsymbol{w}_i \rangle \le c_i(x,z) \qquad \forall (x,z) \in \mathcal{X}_i \times \mathcal{Z},\ \forall 1 \le i \le N, \\ & \sum_{i=1}^N \boldsymbol{w}_i = \boldsymbol{0}. \end{aligned} \qquad (\mathrm{MT}^*_{\mathrm{par}})$$

Duality results

- $(\mathrm{MT}^*_{\mathrm{par}})$ is a **linear semi-infinite programming (LSIP) problem** and admits the following **dual optimization problem**:

$$\begin{aligned} \underset{(\theta_i)}{\text{minimize}} \quad & \sum_{i=1}^{N} \int_{\mathcal{X}_i \times \mathcal{Z}} c_i \, \mathrm{d}\theta_i \\ \text{subject to} \quad & \theta_i \in \Gamma(\bar{\mu}_i, \bar{\nu}_i) && \forall 1 \le i \le N, \\ & \bar{\mu}_i \overset{\mathcal{G}_i}{\sim} \mu_i && \forall 1 \le i \le N, \\ & \bar{\nu}_i \overset{\mathcal{H}}{\sim} \bar{\nu}_1 && \forall 1 \le i \le N. \end{aligned} \qquad (\mathrm{MT}_{\mathrm{par}})$$

- Strong duality can be established via classical LSIP theory [Goberna and López 1998].

Theorem (Strong duality)

The strong duality between $(\mathrm{MT}^*_{\mathrm{par}})$ and $(\mathrm{MT}_{\mathrm{par}})$ holds, i.e., the optimal values of $(\mathrm{MT}^*_{\mathrm{par}})$ and $(\mathrm{MT}_{\mathrm{par}})$ are identical.

Computational complexity

- The theoretical computational complexity of $(\mathrm{MT}^*_{\mathrm{par}})$ and $(\mathrm{MT}_{\mathrm{par}})$ can be analyzed through the **volumetric center method** of Vaidya [1996] in terms of the **global minimization oracle** defined as follows.

Definition (Global minimization oracle)

$\texttt{Oracle}(i, \boldsymbol{y}_i, \boldsymbol{w}_i)$ solves the global minimization problem:

$$\min_{(x,z)\in\mathcal{X}_i\times\mathcal{Z}} \left\{c_i(x,z) - \langle \boldsymbol{g}_i(x), \boldsymbol{y}_i\rangle - \langle \boldsymbol{h}(z), \boldsymbol{w}_i\rangle\right\}$$

and returns an optimizer $(x^\star, z^\star)$ and the optimal value $\beta^\star$ with computational cost T. (Note that T does not depend on N.)

- **Intuition:** `Oracle` determines the "most violated" constraint.

Computational complexity

Theorem (Computational complexity)

Let $m := \sum_{i=1}^{N} m_i = \sum_{i=1}^{N} |\mathcal{G}_i|$, $k := |\mathcal{H}|$. Then, in the Euclidean case (i.e., when $\mathcal{X}_1, \ldots, \mathcal{X}_N, \mathcal{Z}$ are all Euclidean),

1. *the computational complexity of computing an ϵ-optimizer of $(\mathsf{MT}^*_{\mathsf{par}})$ is $O\big((m+Nk)\log((m+Nk)/\epsilon)(NT+(m+Nk)^{\omega})\big)$;*
2. *the computational complexity of computing a pair of ϵ-optimizers of $(\mathsf{MT}^*_{\mathsf{par}})$ and $(\mathsf{MT}_{\mathsf{par}})$ is polynomial in N, m, k, T, and $\log\left(\frac{1}{\epsilon}\right)$.*

$O(m^{\omega})$ is the computational complexity of the multiplication of two $m \times m$ matrices.

- We also derive the theoretical computational complexity in the general non-Euclidean case.

Construction of approximate matching equilibrium

- **Given:** approx. optimizer $(\hat{\theta}_i)_{i=1:N}$ of (MT_{par}), approx. optimizer $(\hat{y}_{i,0}, \hat{\boldsymbol{y}}_i, \hat{\boldsymbol{w}}_i)_{i=1:N}$ of $(\text{MT}^*_{\text{par}})$.
- **Construction of approximate optimizer of (MT^*):**
  - for $i = 1, \ldots, N-1$, $\tilde{\varphi}_i(z) := \inf_{x \in \mathcal{X}_i} \{c_i(x,z) - \hat{y}_{i,0} - \langle \boldsymbol{g}_i(x), \hat{\boldsymbol{y}}_i \rangle\} - \tilde{\varphi}_{i,0}$,
    where $\tilde{\varphi}_{i,0} := \inf_{x \in \mathcal{X}_i} \{c_i(x,z_0) - \hat{y}_{i,0} - \langle \boldsymbol{g}_i(x), \hat{\boldsymbol{y}}_i \rangle\} - \tilde{\varphi}_{i,0}$ for some point $z_0 \in \mathcal{Z}$;
  - $\tilde{\varphi}_N(z) := -\sum_{i=1}^{N-1} \tilde{\varphi}_i(z)$.
- **Construction of approximate optimizer of (MT) via gluing:**
  - **step 1:** for $i = 1, \ldots, N$, glue $\hat{\theta}_i \in \Gamma(\hat{\mu}_i, \hat{\nu}_i)$ with a W_1-optimal coupling $\hat{\eta}_i \in \Gamma(\hat{\mu}_i, \mu_i)$ and a W_1-optimal coupling $\hat{\rho}_i \in \Gamma(\hat{\nu}_i, \hat{\nu}_1)$ to get $\hat{\gamma}_i \in \Gamma(\mu_i, \hat{\nu}_1)$;
  - **step 2:** glue $\hat{\gamma}_1 \in \Gamma(\mu_1, \hat{\nu}_1), \ldots, \hat{\gamma}_N \in \Gamma(\mu_N, \hat{\nu}_1)$ together to get $\tilde{\mu} \in \Gamma(\mu_1, \ldots, \mu_N)$;
  - **step 3:** let $\tilde{\nu} := \tilde{z} \sharp \tilde{\mu}$, and let $\tilde{\gamma}_i := (\text{proj}_i, \tilde{z}) \sharp \tilde{\mu}$ for $i = 1, \ldots, N$, where $\tilde{z}(x_1, \ldots, x_N) \in \arg\min_{z \in \mathcal{Z}} \left\{ \sum_{i=1}^N c_i(x_i, z) \right\}$.

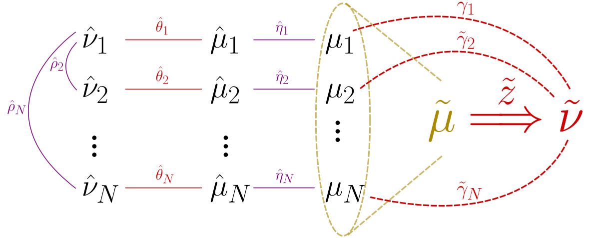

Construction of approximate matching equilibrium

Theorem (Approximate matching equilibrium)

Suppose that:

- c_i *is* L_c*-Lipschitz continuous for* $i = 1, \ldots, N$*;*
- $\epsilon_0 := \left(\sum_{i=1}^{N} \int_{\mathcal{X}_i \times \mathcal{Z}} c_i \,\mathrm{d}\hat{\theta}_i\right) - \left(\sum_{i=1}^{N} \hat{y}_{i,0} + \langle \bar{\boldsymbol{g}}_i, \hat{\boldsymbol{y}}_i \rangle\right)$, *i.e., duality gap when solving* $(\mathrm{MT}_{\mathrm{par}})$ *and* $(\mathrm{MT}^*_{\mathrm{par}})$.
- $\epsilon := \epsilon_0 + L_c \left(\sum_{i=1}^{N} \sup_{\mu_i' \overset{\mathcal{G}_i}{\sim} \mu_i} \left\{W_1(\mu_i, \hat{\mu}_i)\right\} + \sup_{\nu \overset{\mathcal{H}}{\sim} \nu'} \left\{W_1(\nu, \nu')\right\}\right)$.

Then, the constructed $\tilde{\nu}$, $(\tilde{\gamma}_i)_{i=1:N}$, *and* $(\tilde{\varphi}_i)_{i=1:N}$ *satisfy:*

1. $\tilde{\nu}$ *is an* ϵ*-optimizer of* (MT)*;*
2. *for* $i = 1, \ldots, N$, $\tilde{\gamma}_i \in \Gamma(\mu_i, \tilde{\nu})$ *and* $\int_{\mathcal{X}_i \times \mathcal{Z}} c_i \,\mathrm{d}\tilde{\gamma}_i \le W_{c_i}(\mu_i, \tilde{\nu}) + \epsilon$.
3. $(\tilde{\varphi}_i)_{i=1:N}$ *is an* ϵ*-optimizer of* (MT^*)*;*

Such $\left((\tilde{\varphi}_i)_{i=1:N}, \tilde{\nu}, (\tilde{\gamma}_i)_{i=1:N}\right)$ *is called an* ϵ***-approximate matching equilibrium****.*

Convergence to true matching equilibrium

Corollary (Convergence to true matching equilibrium)

Let $((\tilde{\varphi}_i^{(l)})_{i=1:N}, \tilde{\nu}^{(l)}, (\tilde{\gamma}_i^{(l)})_{i=1:N})$ *be an* $\epsilon^{(l)}$*-approx. matching equilibrium with* $\lim_{l\to\infty} \epsilon^{(l)} = 0$*. Then:*

- $(\tilde{\varphi}_i^{(l)})_{l\in\mathbb{N}}$ *has at least one accumulation point in* $(\mathcal{C}(\mathcal{Z}), \|\cdot\|_\infty)$ *for* $i = 1, \ldots, N$*;*
- $(\tilde{\nu}^{(l)})_{l\in\mathbb{N}}$ *has at least one accumulation point in* $(\mathcal{P}(\mathcal{Z}), W_1)$*;*
- $(\tilde{\gamma}_i^{(l)})_{l\in\mathbb{N}}$ *has at least one accumulation point in* $(\mathcal{P}(\mathcal{X}_i \times \mathcal{Z}), W_1)$ *for* $i = 1, \ldots, N$*.*

If $\tilde{\varphi}_i^{(l_t)} \xrightarrow[t\to\infty]{\text{unif.}} \tilde{\varphi}_i^{(\infty)}\ \forall 1 \le i \le N$*,* $\tilde{\nu}^{(l_t)} \xrightarrow[t\to\infty]{W_1} \tilde{\nu}^{(\infty)}$*, and* $\tilde{\gamma}_i^{(l_t)} \xrightarrow[t\to\infty]{W_1} \tilde{\gamma}_i^{(\infty)}\ \forall 1 \le i \le N$*,*
then $((\tilde{\varphi}_i^{(\infty)})_{i=1:N}, \tilde{\nu}^{(\infty)}, (\tilde{\gamma}_i^{(\infty)})_{i=1:N})$ *is a matching equilibrium.*

- In the Euclidean case, we can explicitly construct continuous piece-wise affine basis functions $\mathcal{G}_1, \ldots, \mathcal{G}_N, \mathcal{H}$ to control the approximation error ϵ to be arbitrarily close to 0.

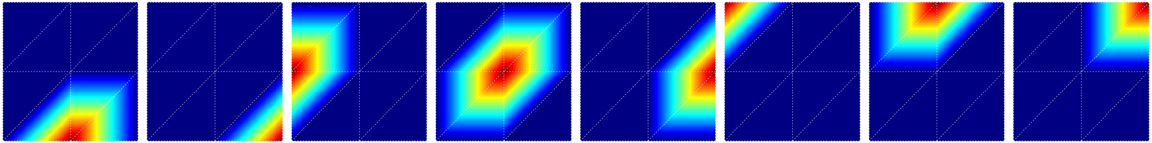

Numerical algorithm

- We first develop a **cutting-plane algorithm** for computing an approximate optimizer $(\hat{\theta}_i)_{i=1:N}$ of $(\mathsf{MT}_{\mathsf{par}})$ and an approximate optimizer $(\hat{y}_{i,0}, \hat{\boldsymbol{y}}_i, \hat{\boldsymbol{w}}_i)_{i=1:N}$ of $(\mathsf{MT}^*_{\mathsf{par}})$.
- We then develop an **algorithm for computing an approximate matching equilibrium** via constructing random variables on a probability space, with the following properties.

Theorem (Matching for teams algorithm)

Under suitable conditions, for any $\epsilon > 0$, the proposed algorithm produces outputs $(\tilde{\varphi}_i)_{i=1:N}$, $\tilde{\nu}$, $(\tilde{\gamma}_i)_{i=1:N}$, α^{LB}, α^{UB}, and ϵ_{sub} satisfying:

1. $\alpha^{\mathsf{LB}} \le (\mathsf{MT}^*) = (\mathsf{MT}) \le \alpha^{\mathsf{UB}}$ *and* $\epsilon_{\mathsf{sub}} := \alpha^{\mathsf{UB}} - \alpha^{\mathsf{LB}} \le \epsilon$ *(typically $\epsilon_{\mathsf{sub}} \ll \epsilon$ in practice);*
2. $\tilde{\nu}$ *is an* ϵ_{sub}*-optimizer of* (MT)*;*
3. *for* $i = 1, \ldots, N$, $\tilde{\gamma}_i \in \Gamma(\mu_i, \tilde{\nu})$ *and* $\int_{\mathcal{X}_i \times \mathcal{Z}} c_i \, \mathrm{d}\tilde{\gamma}_i \le W_{c_i}(\mu_i, \tilde{\nu}) + \epsilon_{\mathsf{sub}}$.
4. $(\tilde{\varphi}_i)_{i=1:N}$ *is an* ϵ_{sub}*-optimizer of* (MT^*)*;*

In particular, $\big((\tilde{\varphi}_i)_{i=1:N}, \tilde{\nu}, (\tilde{\gamma}_i)_{i=1:N}\big)$ *form an* ϵ_{sub}***-approximate matching equilibrium**.*

Experiment 1: business locations

- **Settings:**
  - $N = 5$; the city has a railway line with 5 stations;
  - **commuting costs:**

$$c_i(\boldsymbol{x},\boldsymbol{z}) := \min\Big\{ \underbrace{\|\boldsymbol{x}-\boldsymbol{z}\|_1}_{\text{walk from home to workplace}}, \min_{1\le j,k\le 5}\big\{ \underbrace{\|\boldsymbol{x}-\boldsymbol{u}_j\|_1}_{\text{walk from home to station } j} + \underbrace{C_{j,k}}_{\text{take train from station } j \text{ to station } k} + \underbrace{\|\boldsymbol{u}_k-\boldsymbol{z}\|_1}_{\text{walk from station } k \text{ to workplace}} \big\}\Big\};$$

  - **restocking cost:** $c_N(\boldsymbol{x},\boldsymbol{z}) := \|\boldsymbol{x}-\boldsymbol{z}\|_1$.

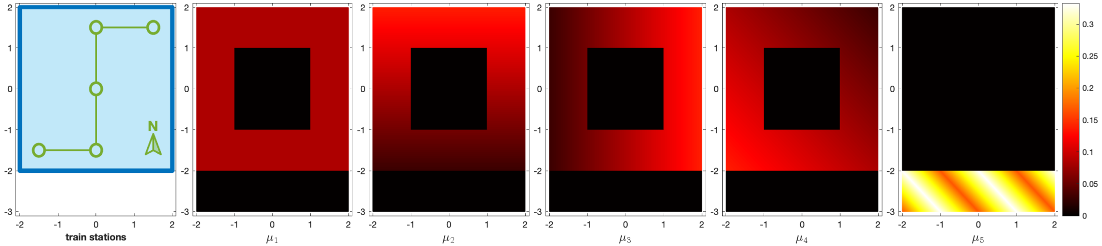

Experiment 1: business locations

- **Computed bounds α^{LB}, α^{UB} and sub-optimality estimates ϵ_{sub}:**

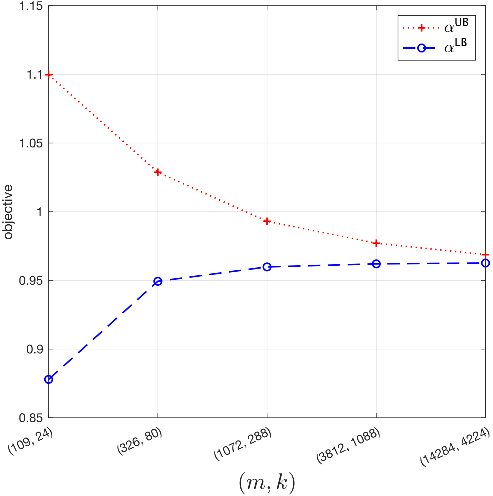

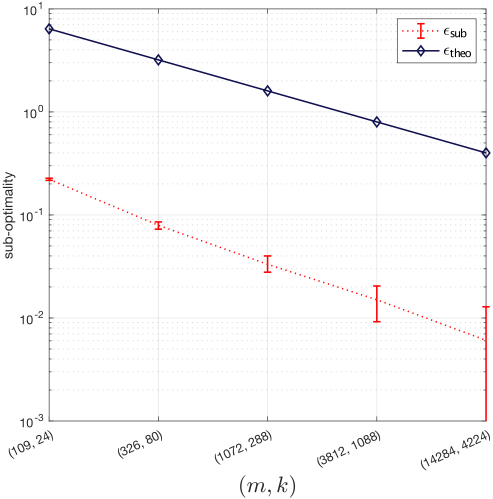

Numerical results

- **Computed approximate matching equilibrium:**

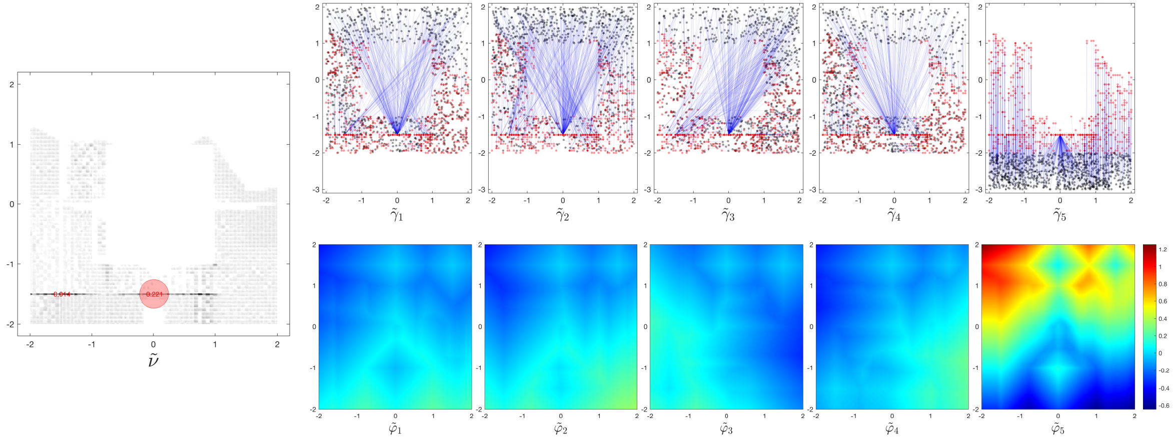

Experiment 2: Wasserstein barycenter

- **Settings:**
  - We compute the Wasserstein barycenter of $N = 20$ probability measures.

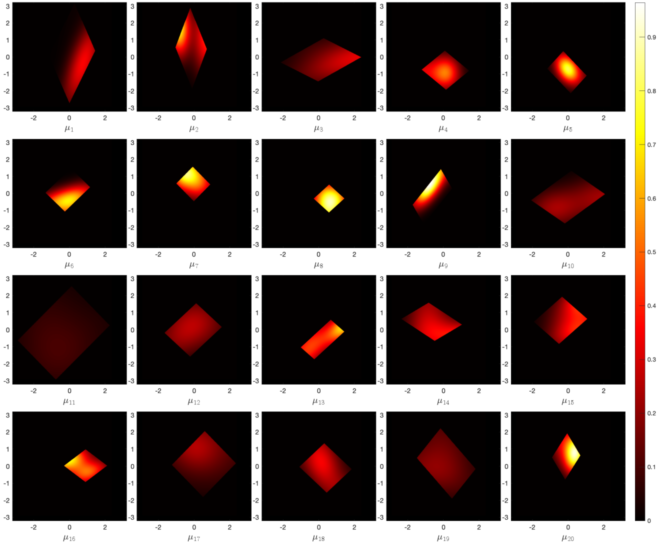

Experiment 2: Wasserstein barycenter

- **Computed bounds α^{LB}, α^{UB} and sub-optimality estimates ϵ_{sub}:**

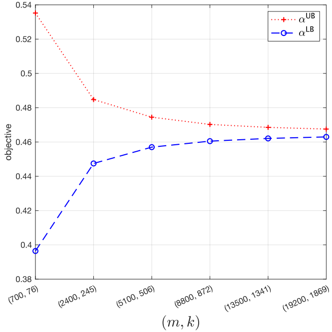

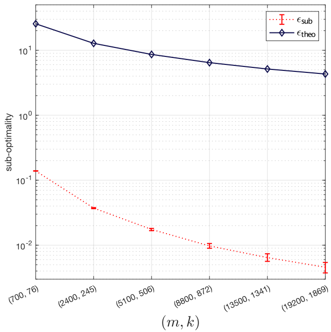

- **Computed approximate Wasserstein barycenters:**

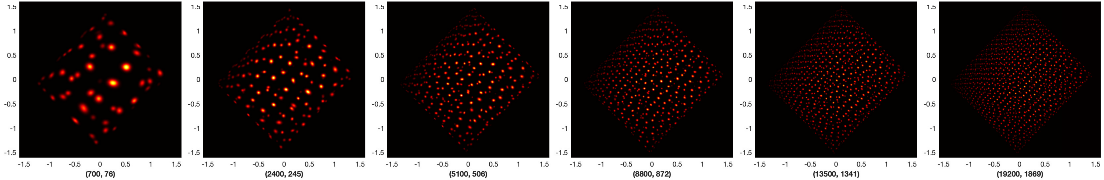

Conclusion

- **Theoretical contributions:**
  - Development of a parametric approximation scheme $(\text{MT}^*_{\text{par}})$ of (MT^*).
  - Derivation of duality results for $(\text{MT}^*_{\text{par}})$ and its dual optimization problem (MT_{par}).
  - Analysis of the theoretical computational complexity of $(\text{MT}^*_{\text{par}})$ and (MT_{par}).
  - Construction of ϵ-optimizers of (MT) and (MT^*) (referred to as ϵ-approximate matching equilibrium), and showing their convergence to a true matching equilibrium.
  - Explicit construction of $(\text{MT}^*_{\text{par}})$ to control the approximation error.
- **Numerical method:**
  - Development of a numerical algorithm which can compute ϵ_0-optimizers of (MT_{par}) and $(\text{MT}^*_{\text{par}})$ for any $\epsilon_0 > 0$.
  - Development of a numerical algorithm which can compute an ϵ-approximate matching equilibrium as well as lower and upper bounds $\alpha^{\text{LB}} \le (\text{MT}^*) = (\text{MT}) \le \alpha^{\text{UB}}$ with $\alpha^{\text{UB}} - \alpha^{\text{LB}} \le \epsilon$ for any $\epsilon > 0$.
  - Application to the business location distribution problem and the Wasserstein barycenter problem.

References

1. **A. Neufeld and Q. Xiang. Feasible approximation of matching equilibria for large-scale matching for teams problems.** *Preprint, arXiv:2308.03550*, 2023.
URL: https://arxiv.org/abs/2308.03550

2. G. Carlier and I. Ekeland. Matching for teams. *Economic Theory*, 42(2):397–418, 2010.

3. M. Agueh and G. Carlier. Barycenters in the Wasserstein space. *SIAM J. Math. Anal.*, 43(2): 904–924, 2011.

4. M. A. Goberna and M. A. López. *Linear semi-infinite optimization.* John Wiley & Sons, 1998.

5. P. M. Vaidya. A new algorithm for minimizing convex functions over convex sets. *Math. Program.*, 73(3):291–341, 1996.